Adaptive Approximations for High-Dimensional Uncertainty Quantification in Stochastic Parametric Electromagnetic Field Simulations

Zur Erlangung des akademischen Grades Doktor-Ingenieur (Dr.-Ing.)
genehmigte Dissertation von Dimitrios Loukrezis aus Athen, Griechenland
Tag der Einreichung: 29. October 2018, Tag der Prüfung: 04. February 2019
Darmstadt — D 17

1. Gutachten: Prof. Dr.-Ing. Herbert De Gersem
2. Gutachten: Prof. Dr.-Ing. Ulrich Römer

TECHNISCHE
UNIVERSITÄT
DARMSTADT

Fachbereich Elektrotechnik
und Informationstechnik
Institut für Theorie
Elektromagnetischer Felder

Adaptive Approximations for High-Dimensional Uncertainty Quantification in Stochastic Parametric Electromagnetic Field Simulations

Genehmigte Dissertation von Dimitrios Loukrezis aus Athen, Griechenland

1. Gutachten: Prof. Dr.-Ing. Herbert De Gersem
2. Gutachten: Prof. Dr.-Ing. Ulrich Römer

Tag der Einreichung: 29. October 2018
Tag der Prüfung: 04. February 2019

Darmstadt — D 17

Bitte zitieren Sie dieses Dokument als:
URN: urn:nbn:de:tuda-tuprints-84854
URL: http://tuprints.ulb.tu-darmstadt.de/84854

Dieses Dokument wird bereitgestellt von tuprints,
E-Publishing-Service der TU Darmstadt
http://tuprints.ulb.tu-darmstadt.de
tuprints@ulb.tu-darmstadt.de

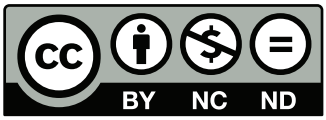

Erklärung zur Dissertation

Hiermit versichere ich, die vorliegende Dissertation ohne Hilfe Dritter nur mit den angegebenen Quellen und Hilfsmitteln angefertigt zu haben. Alle Stellen, die aus Quellen entnommen wurden, sind als solche kenntlich gemacht. Diese Arbeit hat in gleicher oder ähnlicher Form noch keiner Prüfungsbehörde vorgelegen.

Darmstadt, den 04. February 2019

_____

(Dimitrios Loukrezis)

Contents

Abstract

The present work addresses the problems of high-dimensional approximation and uncertainty quantification in the context of electromagnetic field simulations. Such problems are typically encountered during the design phase of electromagnetic devices, e.g. magnets or high-frequency components. Manufacturing tolerances, material contaminations, or other types of imperfections introduce, possibly many, sources of uncertainty with respect to the device's characteristics, e.g. geometry, material properties, or operational data, which in turn affect the overall operation of the device. For the final designs to be robust and the manufactured device to operate within its specifications, this uncertainty must be accounted for in the simulation-based studies performed during the design phase.

In the presence of many parameters, one faces the so-called curse of dimensionality, i.e. the computational complexity increases exponentially as the number of parameters, equivalently, dimensions, grows. The focus of this work lies on adaptive methods that mitigate the effect of the curse of dimensionality, and therefore enable otherwise intractable uncertainty quantification studies. Its application scope includes electromagnetic field models suffering from moderately high-dimensional input uncertainty. However, the presented methods can be used in a black-box fashion and are therefore applicable to other types of problems as well, such as fluid dynamics or structural mechanics, provided that the underlying assumptions, e.g. smoothness, are satisfied. To that end, three different approaches are investigated.

The first approach relies on a dimension-adaptive stochastic collocation scheme. Emphasis is placed on the use of the commonly employed Clenshaw-Curtis collocation points and the relatively recently investigated Leja collocation points, for both uniform and non-uniform input distributions. It is shown that combining the stochastic collocation method with a well-known dimension-adaptive algorithm results in significant computational savings compared to the isotropic collocation variant. For the case of uniformly distributed input parameters, the performance of Leja nodes is found to be advantageous in terms of approximation accuracy, but inferior for quadrature purposes, compared to Clenshaw-Curtis nodes. The reverse is observed for the case of skewed beta input distributions, where Leja nodes yield inferior approximation accuracies, but offer significant advantages in the quadrature context.

The second approach employs adaptively constructed generalized polynomial chaos expansions. A two-level adaptivity is considered. The first adaptivity level refers to the construction of the polynomial basis given an experimental design of fixed size. The second adaptivity level refers to the adaptive expansion of the experimental design, whenever necessary due to numerical stability requirements. We note that a drawback of this second level of adaptivity is its dependence on an

a priori set limit value $\kappa_{\max}$ regarding the condition number of a system matrix, which acts as a stability indicator. Based on this two-level adaptivity approach, an in-house algorithm has been developed in the course of this work. The algorithm is presented in detail and tested against isotropic polynomial chaos expansions, showing significant computational savings. The algorithm is also compared against a state-of-the-art adaptive polynomial chaos approach and is found to be superior, even for conservative, pessimistic values of $\kappa_{\max}$.

The third approach is based on low-rank tensor decompositions. Exploiting the underlying tensor structure of the multivariate quadrature formulas which are employed for the numerical computation of statistical moments, tensor decompositions are employed to reduce the complexity of the multi-dimensional arrays, i.e. tensors. First, it is shown how quadrature weight tensors may admit exact low-rank representations. Next, high-order adaptive cross approximation methods are used to compute low-rank approximations of the function-generated tensors which contain the model evaluations on the quadrature nodes. The resulting low-rank tensor decompositions are then used for the estimation of statistical moments at a reduced cost. The numerical results show that the approach yields tremendous complexity reductions compared to full tensor-product quadratures, however, it is significantly more expensive compared to the aforementioned collocation and polynomial chaos approaches.

The implementations of all considered approaches are documented in in-house developed Python and MATLAB scripts and software libraries. Extensive numerical tests are presented for all methods with respect to their accuracy and computational cost. The methods are tested on simulation models of both toy examples and real-world electromagnetic field applications, featuring a moderately high number of uncertain parameters.

Kurzfassung

Die vorliegende Arbeit befasst sich mit hochdimensionaler Approximation und Unsicherheitsquantifizierung im Kontext elektromagnetischer Feldsimulation. Solche Probleme treten typischerweise während der Entwurfsphase eines elektromagnetischen Geräts auf, z.B. eines Magneten oder einer Hochfrequenzkomponente. Herstellungstoleranzen, Materialkontamination oder andere Mängel stellen Unsicherheitsquellen in Bezug auf die Eigenschaften des Gerätes dar, die dessen Bedienung beeinflussen können. Um eine robuste und zuverlässige Funktionsweise, am Ende des Entwurfsprozesses sicherszustellen, müssen Unsicherheiten in simulationsbasierten Analysen, die während der Entwurfsphase durchgeführt werden, berücksichtigt werden.

Wenn Unsicherheiten in vielen Eingangsparametern vorliegen, tritt der sogenannte Fluch der Dimensionalität auf, d.h. die Komplexität steigt exponentiell mit der Anzahl von Parametern. Der Schwerpunkt dieser Arbeit liegt auf der Verwendung von adaptiven Methoden, die den Effekt des Fluches der Dimensionalität abschwächen und daher aufwendige Unsicherheitsquantifizierungsstudien ermöglichen. Der Anwendungsbereich umfasst elektromagnetische Feldmodelle, die eine moderat hochdimensionale Eingangsunsicherheit aufweisen. Die vorgestellten Verfahren sind jedoch auch auf andere Arten von Problemen anwendbar, z.B. aus den Bereichen der Fluiddynamik oder Strukturmechanik. Zu diesem Zweck werden drei verschiedene Ansätze untersucht.

Der erste Ansatz basiert auf einer dimensionsadaptiven stochastischen Kollokationsmethode. Der Schwerpunkt liegt auf der Verwendung von Leja und Clenshaw-Curtis Kollokationspunkten für gleichverteilte und nicht gleichverteilte Zufallsvariabeln. Es wird gezeigt, dass die Kombination der stochastischen Kollokationsmethode mit einem bekannten dimensionsadaptiven Algorithmus zu signifikanten rechnerischen Einsparungen im Vergleich zur isotropen Kollokationsvariante führt. Für den Fall gleichverteilter Parameter haben sich Leja-Knoten im Vergleich zu Clenshaw-Curtis-Knoten, im Hinblick auf die Genauigkeit der Approximation, als vorteilhaft erwiesen, für Quadraturzwecke jedoch als unterlegen. Das Umgekehrte wird für den Fall von nicht-symmetrischen Betaverteilungen beobachtet, wo Leja-Knoten geringere Approximationsgenauigkeiten ergeben, jedoch signifikante Vorteile im Quadraturkontext bieten.

Der zweite Ansatz verwendet adaptive generalisierte Polynomiale-Chaos Entwicklungen. Dabei wird eine zweistufige Adaptivität verwendet. Die erste Adaptivitätsstufe bezieht sich auf die Konstruktion der Polynombasis bei einer Versuchsplanung fester Größe. Die zweite Adaptivitätsstufe bezieht sich auf die adaptive Erweiterung der Versuchsplanung, aufgrund numerischer Stabilitätsanforderungen. Basierend auf diesem zweistufigen Adaptivitätsansatz wurde im Rahmen dieser

Arbeit ein Algorithmus entwickelt. Der Algorithmus wird im Detail vorgestellt und gegen isotrope polynomiale Chaosexpansionen getestet, wobei signifikante Einsparungen bei der Rechenzeit erzielt werden. Der Algorithmus wird darüber hinaus auch mit einem state-of-the-art Polynomialen-Chaos-Ansatz verglichen, dem gegenüber er überlegen ist.

Der dritte Ansatz basiert auf Niedrigrang-Tensorzerlegungen. Dazu wird die zugrundeliegende Tensorstruktur von multivariaten Quadraturformeln genutzt, die für die numerische Berechnung statistischer Momente verwendet werden. Tensorzerlegungen werden angewandt, um die Komplexität der mehrdimensionalen Arrays, d.h. Tensoren, zu reduzieren. Zuerst wird gezeigt, wie Tensoren bestehend aus Quadraturgewichten eine exakte Niedrigrang-Darstellung zulassen. Als Nächstes werden adaptive Kreuznäherungsverfahren höherer Ordnung angewandt, um die funktionserzeugten Tensoren in einem Niedrigrang-Format zu approximieren, die die Modellauswertungen auf den Quadraturknoten enthalten. Die resultierenden Niedrigrang-Tensorzerlegungen werden dann für die Schätzung von statistischen Momenten verwendet. Es wird gezeigt, dass dieser Ansatz im Vergleich zu vollständigen Tensor-Produkt-Quadraturen zu enormen Komplexitätsreduktionen führt. Der Ansatz ist jedoch im Vergleich zu den oben genannten Kollokations- und Polynomialen-Chaos-Ansätzen deutlich teurer.

Die Implementierungen aller Ansätze sind in selbstentwickelten Python- und MATLAB-Skripten sowie Software-Bibliotheken dokumentiert. Umfangreiche numerische Tests werden für alle Methoden hinsichtlich ihrer Genauigkeit und ihres Rechenaufwandes präsentiert. Die Methoden werden an Simulationsmodellen von sowohl stark vereinfachten Simulationsmodellen als auch an realen elektromagnetischen Feldproblemen getestet, die eine moderat hohe Anzahl von unsicheren Parametern aufweisen.

1 Introduction

In this introductory chapter we present general information about the present thesis. The chapter begins with the motivation that led to this particular thesis topic. Next, we discuss the contributions of the present work withing this research field. The structure of the thesis is presented at the end of the chapter.

1.1 Motivation

Various engineering branches employ nowadays "in silico" models [68], i.e. computer programs or softwares which simulate physical phenomena by solving the underlying mathematical problems on a computer. Simulation-based parameter studies, such as optimization, sensitivity, or uncertainty analyses, allow for thorough examinations of the relations between a model's input parameters and its output quantities. The results of such studies are often used to improve the designs of actual physical systems and devices, such as magnets, electrical machines, waveguides or antennas, to name a few indicative examples.

The relatively young field of uncertainty quantification (UQ) has emerged among the computational sciences to address the need of taking uncertainty into account when performing simulation-based studies. Although state-of-the-art models become increasingly accurate thanks to sophisticated computational methods and powerful computer hardware and software, uncertainties with respect to the model parameters gives rise to uncertainties regarding the model output quantities as well. Let us for example consider the case of designing and eventually producing an electromagnetic device. While parameter values which result in the expected operation of the device can be identified during the design phase, manufacturing or other tolerances may cause the actual values of said parameters to deviate from the desired ones. As a consequence, a suboptimal operation or, in extreme cases, even a failure of the device are possible outcomes. Therefore, uncertainty must be taken into consideration in order to minimize the risk of such events.

A common bottleneck in multi-parameter studies is the so-called "curse of dimensionality". The term is attributed to R. E. Bellman [18] and refers to a variety of difficulties which arise in computational tasks performed in high-dimensional spaces. Most commonly, the curse of dimensionality manifests in the form of paralyzing complexities which render the computational costs of standard methods unaffordable. High-dimensional UQ is no stranger to such problems. For example, estimating a statistical moment of a model output typically reduces to computing a multivariate integral. As illustrated in Figure 1.1, a numerical integration scheme based on tensor-product combinations of univariate quadrature rules becomes quickly intractable for an increasing number of parameters. Since the strength of modern simulation models lies, to a large extent, in their capability to

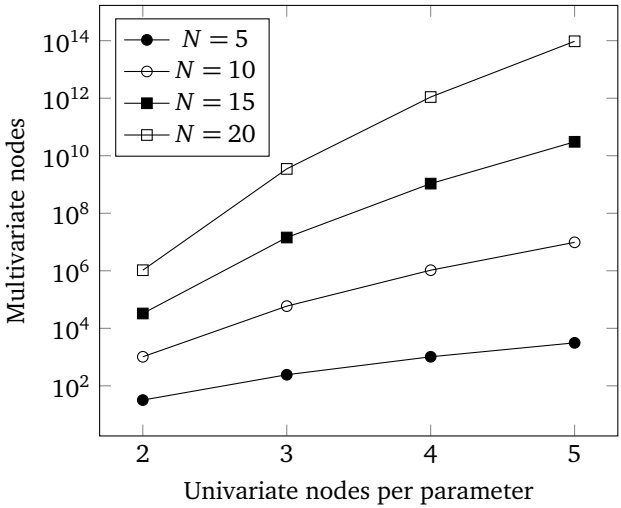

Figure 1.1: Exponential complexity growth of a simple quadrature scheme with N dimensions/parameters.

incorporate a large number of parameter dependencies, the development of methods which allow to mitigate or even break the curse of dimensionality is currently a very active field of research[1].

Methods such as active subspaces [43] and model order reduction [35, 74] attempt a dimensionality reduction before proceeding to the actual parameter study. In the context of stochastic electromagnetic field (EMF) simulations, such methods have been successfully applied in [19, 79, 134], among other works. However, dimensionality reduction is not always possible and in many cases a large number of parameters must be considered as not to compromise the accuracy. The classic Monte Carlo (MC) sampling approach [31, 90] is not affected by the number of parameters, however, its slow convergence rate makes the method unattractive when a high accuracy is required. An improved convergence can be achieved with quasi-MC [31, 90] and multilevel MC [62] methods, under certain assumptions.

Assuming some regularity in the model's input-to-output map, a high-order algebraic, or even exponential convergence can be achieved with the use of UQ methods based on spectral approximations [61, 87, 162]. Spectral UQ methods have

[1] There is a certain irony here, also pointed out in [43], in the sense that increased model complexity is supposed to be a means to reduce uncertainty. Instead, it often becomes a gateway to further uncertainty sources.

been applied to problems from fluid dynamics [87, 147], structural mechanics [22] and microwave engineering [5], to mention just a few application domains. The stochastic Galerkin method [8, 61, 101] is considered to be the optimal one in terms of the accuracy-cost ratio, however, the method requires a dedicated solver for each considered problem. This requirement is usually regarded as a major drawback, especially in the case of complex simulation models where access "under the hood" is not provided, e.g. when a commercial software is used as a black-box. In most practical applications, scientists and engineers turn to the stochastic collocation method [6], or to generalized polynomial chaos (gPC) [164] approximations based on regression [24, 106] or pseudo-spectral projection [49, 88]. Those approaches are preferred because they achieve accuracy-cost ratios which are comparable to the stochastic Galerkin ones, while at the same time being non-intrusive, i.e. they allow a black-box use of the simulation model. We note that there is no universally agreed definition regarding the "intrusiveness" of a method [63], however, here we use the usual distinction.

In high dimensions, adaptive approximations are used to circumvent the curse of dimensionality. Adaptive stochastic collocation approaches have been considered in [37, 60, 83, 96, 110, 117, 139]. Methods and algorithms for adaptive gPC can be found in [2, 22, 24, 102, 113]. More recently, on the basis of the observation that the multivariate polynomial bases employed in spectral approximations have an underlying tensor structure, methods based on tensor decompositions [69, 70, 84] and low-rank tensor approximations [11, 12, 124, 137, 138, 167] have gained increasing interest. The main idea behind all adaptive methods is that, in most practical cases, the output of a model is rarely affected equally by variations from all model parameters, i.e. some parameters are more significant than others. Assuming that such a parameter anisotropy exists, adaptive methods and the corresponding algorithms are employed to separate the significant parameter contributions from the negligible ones. The consequent anisotropic investment of computational resources results in computational savings without compromising the desired accuracy.

The focal point of the present work is to enable high-dimensional UQ in the context of EMF simulations, exactly by exploiting the aforementioned parameter anisotropy. We note that we focus on uncertainty propagation problems only. Other types of UQ studies, such as Bayesian parameter estimation [140] or model-form uncertainty [131], while definitely interesting, are not pursued in this work. We consider adaptive approaches for UQ purposes in the contexts of stochastic collocation, gPC and tensor decompositions. Both established methods and new algorithms, developed as part of this work, are applied to EMF simulation models suffering from moderate to high-dimensional input uncertainty. While we only

consider here applications related to EMF problems, we note that the application scope of all considered approaches is vastly broader.

1.2 Contribution

As indicated by the number of already cited works, the methods which are examined in the present thesis have been extensively analyzed and used in various fields and applications. We consider the main contributions of this work to be the following:

1. The implementation of a dimension-adaptive stochastic collocation algorithm and its application for uniform and fairly "exotic" input probability density functions (PDFs), in particular skewed beta ones. The latter case has not been considered in the literature so far. Moreover, the method has hardly been applied to concrete engineering applications. We have contributed in this direction with our works [57, 96].

2. The development of an algorithm for the adaptive construction of gPC approximations. The adaptive expansion of both the gPC polynomial basis and the experimental design is considered. To the author's knowledge, the proposed adaptive scheme is new and distinctly different from other adaptive gPC algorithms available in the literature, see e.g. [2, 22, 24, 102, 113].

3. The application of tensor decompositions for statistical moment estimations and their comparison against isotropic and adaptive stochastic collocation and gPC methods with respect to accuracy and computational effort. As in the collocation case, complicated engineering examples have barely been considered in the literature. In the context of EMF simulations, we are aware of the work of Zhang et al. [167] and our own contribution [97].

4. The application of all aforementioned approaches to a number of EMF applications, ranging from low- to high-frequency electromagnetics. Both simple toy examples and real-world models have been used for verification purposes.

5. The development of Python and MATLAB software libraries for all of the aforementioned methods and algorithms.

1.3 Structure

The present thesis is organized as follows. Chapter 2 discusses some necessary preliminary notions which will be used throughout the thesis. The next three chapters present the methods and approaches which are examined in this work. In Chapter 3 we recall the basics of the stochastic collocation method and present a

hierarchical collocation scheme based on nested collocation points. Chapter 4 first presents the theory behind gPC approximations and then discusses a number of algorithms for its adaptive construction, in terms of both basis and experimental design expansion. Chapter 5 presents the notion of tensors and tensor decompositions in the context of multi-linear algebra. Common tensor decompositions are presented, along with a discussion on their strengths and weaknesses. Emphasis is placed on available black-box methods and algorithms for the adaptive cross approximation of tensors and their use in the context of UQ. In all three methodology-related chapters, an academic waveguide model with random input data is used to verify the advantages of the proposed adaptive methods. The penultimate Chapter 6 presents the application of all aforementioned approaches to moderately high-dimensional stochastic parametric models. Both analytical and discretized models with $11 - 14$ uncertain parameters are considered, the latter based on the finite element method (FEM) or on isogeometric analysis (IGA). The thesis ends with a conclusion on the findings of the presented work and with a discussion on possible extensions, both available in Chapter 7.

2 Preliminaries

In this chapter, we discuss some necessary preliminaries which will be used throughout this thesis. The chapter starts by providing brief explanations of the terms “parametric” and “stochastic parametric”, which characterize the here-considered problems and models. For further clarification, a stochastic parametric EMF problem and its mathematical model are presented, namely a dielectric slab waveguide with random inputs. We proceed with the description of the objectives of uncertainty propagation and quantification. In the same section we provide the definitions of the error metrics which will be used in subsequent chapters. Next, we offer a short introduction to spectral UQ methods. The chapter continues with definitions related to the downward-closedness property of a set, which is often recalled in subsequent chapters. Finally, general notions on adaptivity are introduced, with emphasis on a posteriori error indicators.

2.1 Stochastic Parametric Models

We start by defining a mathematical model as the collection of variables and the relations among them, which are used in order to describe a phenomenon in the form of equations [68, 119]. In the context of this work, a mathematical model shall always describe a physical phenomenon, in particular related to EMFs. However, the use of mathematical models extends well beyond the natural sciences. The digital form of a mathematical model, enabling the study of the underlying physical problem on a computer, will be referred to as the computational, or computerized, or in-silico, or simulation model [68, 119]. We distinguish between a simulation model and a simulation, such that the latter refers to a model evaluation for a specific model configuration, e.g. regarding its solver settings or input values. Evaluating a model is often encountered in the literature as “running” a simulation or “calling” a model.

In this work we employ parametrized simulation models, the predictions of which depend on a set of input parameters. We only consider parameters which affect the underlying mathematical model, e.g. the size and shape of the computational domain or certain material properties. Parameters related to discretization, tolerances or other solver-specific settings of the simulation model are assumed to be chosen accurately enough for the considered application and are not considered as model input parameters. The goal of parameter studies is to use the aforementioned parametrized models in order to examine the relation between their input parameters and one or several model outputs, commonly referred to as the quantities of interest (QoIs).

Table 2.1: Notation summary for parametric problems.

Symbol	Explanation
$\mathbf{y} \in \mathbb{R}^N$	N-dimensional parameter vector
$g : \mathbb{R}^N \to \mathbb{R}$	map from the input parameters to the QoI
$G(u(\mathbf{y})) = g(\mathbf{y})$	parameter-dependent QoI

Let us attempt a first formalization of the parametric problem setting. We assume that the physical problem at hand is tackled by solving numerically a set of parametric partial differential equations (PDEs), given in the general form

$$\mathscr{D}(u, \mathbf{y}) = 0, \tag{2.1}$$

where $\mathscr{D} = \mathscr{D}(u, \mathbf{y})$ is a differential operator, $u = u(\mathbf{y})$ is the solution of (2.1) and $\mathbf{y} \in \mathbb{R}^N$ is an N-dimensional parameter vector. The QoI is typically given as a functional of the solution, here denoted with $G(u(\mathbf{y}))$. For simplicity, we assume that $G(u(\mathbf{y})) \in \mathbb{R}$, however, complex and/or vector-valued QoIs may also be considered. We denote the map from the input parameters to the QoI with $g : \mathbf{y} \mapsto g(\mathbf{y})$, where $g(\mathbf{y}) = G(u(\mathbf{y}))$. A summary of the notation is presented in Table 2.1. The model itself is assumed to be deterministic; the exact same result $g(\mathbf{y})$ is produced each and every time a simulation is run for the same parameter vector $\mathbf{y}$.

Stochasticity enters the problem setting in the form of uncertain (random) input parameters, the values of which are not exactly known a priori and vary randomly in a defined range. The random input parameters are typically modeled as an N-dimensional multivariate random variable (RV) $\mathbf{Y} = (Y_1, Y_2, \ldots, Y_N)$, also called a random vector. The random vector $\mathbf{Y}$ is defined on a complete probability space (Θ, Σ, P), where Θ denotes the sample (outcome) space, Σ the σ-algebra (set of events) and P the probability measure, i.e. a map $P : \Sigma \to [0, 1]$ from events to probabilities. The single RVs Y_n, $n = 1, 2, \ldots, N$, are functions from the corresponding outcome spaces to the measurable spaces $\Xi_n \subseteq \mathbb{R}$, called the image spaces, and are characterized by the univariate PDFs $\varrho_n(y_n)$, such that $\varrho_n : \Xi_n \to \mathbb{R}_{\geq 0}$. Denoting the multidimensional image space with $\Xi \subseteq \mathbb{R}^N$, the random vector $\mathbf{Y}$ can be similarly defined as the map $\mathbf{Y} : \Theta \to \Xi$ and is characterized by the joint PDF $\varrho(\mathbf{y})$, such that $\varrho : \Xi \to \mathbb{R}_{\geq 0}$. Then, the parameter vector represents a realization of the random vector, such that $\mathbf{y} = \mathbf{Y}(\theta) \in \Xi$, $\theta \in \Theta$. Assuming that the random vector

Table 2.2: Notation summary for stochastic parametric problems.

Symbol	Explanation
(Θ, Σ, P)	complete probability space
Θ	sample/outcome space
Σ	σ-algebra/set of events
$P : \Sigma \to [0,1]$	probability measure
Ξ	image set
$\mathbf{Y} : \Theta \to \Xi$	N-dimensional random vector
$\varrho : \Xi \to \mathbb{R}_{\geq 0}$	probability density function (PDF)
$\mathbf{y} = \mathbf{Y}(\theta)$	random realization

consists of mutually independent RVs, the multidimensional image set Ξ and the joint PDF $\varrho(\mathbf{y})$ are given as

$$\Xi = \Xi_1 \times \Xi_2 \times \cdots \times \Xi_N, \tag{2.2}$$

$$\varrho(\mathbf{y}) = \prod_{n=1}^{N} \varrho_n(y_n). \tag{2.3}$$

Table 2.2 offers a summary of the notation.

The deterministic map g of the now random input parameters results in a random output $g(\mathbf{Y})$. In other words, the input uncertainty propagates through the deterministic model and renders the QoI uncertain as well[2]. Problems of this type fall into the category of forward UQ or uncertainty propagation, which is the topic of Section 2.2. Before proceeding to the presentation of the related UQ concepts, we first clarify the here-presented stochastic parametric setting with a model problem from the field of high frequency (HF) electromagnetics.

2.1.1 Model Problem: Dielectric Slab Waveguide with Random Inputs

The content of this section is partially based on our contribution [96, Section 4.1]. We consider a three-dimensional, rectangular dielectric slab waveguide, as the one illustrated in Figure 2.1. As can be seen from Figure 2.1a, the computational domain has a dielectric material with permittivity $\varepsilon = \varepsilon_0 \varepsilon_\mathrm{r}$ and permeability $\mu = \mu_0 \mu_\mathrm{r}$ in its middle (yellow area), while the rest is filled with vacuum (blue areas). The subscripts "0" and "r" refer to the absolute value of the material property in vacuum and its relative value for the given dielectric material, respectively. The red planes denote the waveguide's input and output ports, respectively port 1

[2] Uncertainty propagation assumes that the Doob–Dynkin lemma [130, Proposition 3] is satisfied.

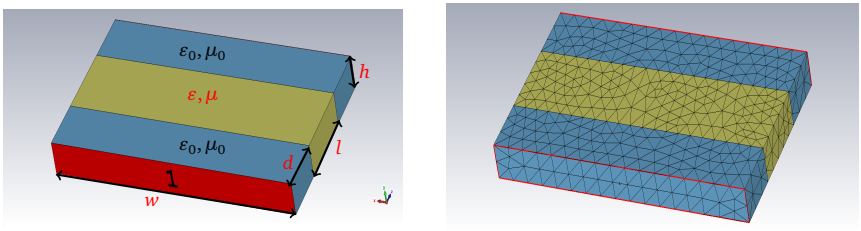

(a) Geometrical and material parameters, denoted with the red script.

(b) Finite element tetrahedral mesh.

Figure 2.1: 3D dielectric slab waveguide model, generated with CST Microwave Studio [44]. The dielectric filling is shown in yellow. The blue areas are vacuum. The red plane with the writing "1" is the input port. The output port "2" lies on the opposite side.

and port 2. Simple waveguide models as this one are typically used to study wave confinement mechanisms [120].

The geometry of the waveguide is defined by its width w, height h, dielectric slab length l and vacuum offset d. All 4 geometrical parameters are depicted in Figure 2.1a. With the exception of the 2 ports, the walls of the waveguide are considered to be perfect electric conductors (PECs). The computational domain Ω and its boundary $\partial\Omega = \Gamma_{\text{PEC}} \cup \Gamma_{\text{in}} \cup \Gamma_{\text{out}}$ are formally defined in the Cartesian coordinate system as

$$\Omega = [0, w] \times [0, h] \times [0, 2d + l], \tag{2.4a}$$
$$\Gamma_{\text{PEC}} = \{(x, y, z) \in \partial\Omega \,:\, z \neq 0 \wedge z \neq 2d + l\}, \tag{2.4b}$$
$$\Gamma_{\text{in}} = \{(x, y, z) \in \partial\Omega \,:\, z = 0\}, \tag{2.4c}$$
$$\Gamma_{\text{out}} = \{(x, y, z) \in \partial\Omega \,:\, z = 2d + l\}. \tag{2.4d}$$

Assuming that the waveguide is excited at port 1 by an incoming wave $\mathbf{U}^{\text{inc}}$, the mathematical model is given by the formulation of Maxwell's source problem for the electric field $\mathbf{E}$. In the following, we will assume that the incoming field coincides with the fundamental transverse electric (TE) mode TE_{10} and that higher

order modes are quickly attenuated in the structure. In this case, Maxwell's source problem reads

$$\begin{aligned}
\operatorname{curl}\left(\mu^{-1}\operatorname{curl}\mathbf{E}\right) - \omega^2\varepsilon\mathbf{E} &= 0, && \text{in } \Omega, && \text{(2.5a)}\\
\mathbf{E}\times\mathbf{n} &= 0, && \text{on } \Gamma_{\text{PEC}}, && \text{(2.5b)}\\
\mathbf{n}\times\operatorname{curl}\mathbf{E} + \gamma\mathbf{n}\times(\mathbf{n}\times\mathbf{E}) &= \mathbf{U}^{\text{inc}}, && \text{on } \Gamma_{\text{in}}, && \text{(2.5c)}\\
\mathbf{n}\times\operatorname{curl}\mathbf{E} + \gamma\mathbf{n}\times(\mathbf{n}\times\mathbf{E}) &= 0, && \text{on } \Gamma_{\text{out}}, && \text{(2.5d)}
\end{aligned}$$

where $\mathbf{n}$ is the outwards-pointing normal vector, $\gamma = \jmath k^{\text{inc}}$, and k^{inc} refers to the wavenumber of the incoming wave $\mathbf{U}^{\text{inc}}$, as in [82].

Typical QoIs for waveguide devices and models are the so-called scattering parameters, S-parameters for short. The S-parameters quantify the reflection and transmission of the incoming field at the ports of the waveguide. For example, the S_{11} parameter, also referred to as the reflection coefficient, quantifies the reflection at port 1 and is given by

$$S_{11} = C^{\text{inc}} \int_{\Gamma_{\text{in}}} \mathbf{E}\cdot\mathbf{e}_{10}\,\mathrm{d}x, \tag{2.6}$$

where C^{inc} is a normalization constant and $\mathbf{e}_{10} = \mathbf{e}_y \sin\frac{\pi x}{w}$ [82], with $\mathbf{e}_y := (0,1,0)$ being the unit vector in the Cartesian y-direction. The S_{11} parameter can take complex values. The QoI is usually chosen to be either $g(\mathbf{y}) = S_{11}(\mathbf{y}) \in \mathbb{C}$, or $g(\mathbf{y}) = |S_{11}(\mathbf{y})| \in [0,1]$.

For most waveguide devices and structures, the unknown electric field $\mathbf{E}$ in problem (2.5) is computed numerically, e.g. the computational domain is approximated by a tetrahedral mesh as in Figure 2.1b, and an approximation $\mathbf{E}_h \approx \mathbf{E}$ is computed with the finite element method (FEM). Then, the scattering parameter S_{11} is obtained by post-processing the discrete solution $\mathbf{E}_h$. However, for this particular, simple waveguide model, an analytical solution for the S_{11} parameter exists and can be used to avoid the consideration of discretization errors. The analytical solution for the dielectric slab waveguide is presented in Appendix A. The FEM formulation for Maxwell's source problem can be found in Section 6.3.1, with respect to a waveguide filter model, for which no analytical solution exists.

We now assume that the geometrical parameters w, h, l, and d, as well as the material parameters ε_{r} and μ_{r} are independent RVs defined on the probability space (Θ, Σ, P) and collect them in the random vector $\mathbf{Y} = (w, h, l, d, \varepsilon_{\text{r}}, \mu_{\text{r}})$. For a given

Table 2.3: Nominal parameter values for the dielectric slab waveguide.

Parameter	Symbol	Nominal Value	Units
width	w	30	mm
height	h	3	mm
filling length	l	7	mm
vacuum offset	d	5	mm
relative permittivity	ε_r	2.0	–
relative permeability	μ_r	2.4	–

realization $\mathbf{y} = \mathbf{Y}(\theta)$, the parametric counterpart of the mathematical model (2.5) reads

$$\begin{aligned}
\operatorname{curl}\left(\mu(\mathbf{y})^{-1}\operatorname{curl}\mathbf{E}\right) - \omega^2\varepsilon(\mathbf{y})\mathbf{E} &= 0, && \text{in } \Omega(\mathbf{y}), && \text{(2.7a)}\\
\mathbf{E}\times\mathbf{n} &= 0, && \text{on } \Gamma_{\text{PEC}}(\mathbf{y}), && \text{(2.7b)}\\
\mathbf{n}\times\operatorname{curl}\mathbf{E} + \gamma\mathbf{n}\times(\mathbf{n}\times\mathbf{E}) &= \mathbf{U}^{\text{inc}}, && \text{on } \Gamma_{\text{in}}(\mathbf{y}), && \text{(2.7c)}\\
\mathbf{n}\times\operatorname{curl}\mathbf{E} + \gamma\mathbf{n}\times(\mathbf{n}\times\mathbf{E}) &= 0, && \text{on } \Gamma_{\text{out}}(\mathbf{y}). && \text{(2.7d)}
\end{aligned}$$

The solution of (2.7) is also parameter-dependent, i.e. $\mathbf{E} = \mathbf{E}(\mathbf{y})$, and the parametric QoI is given by

$$S_{11}(\mathbf{y}) = C^{\text{inc}} \int_{\Gamma_{\text{in}}(\mathbf{y})} \mathbf{E}(\mathbf{y})\cdot\mathbf{e}_{10}\,\mathrm{d}x. \tag{2.8}$$

A common approach in the context of UQ with random geometries is to pull back the parametric equations to a fixed reference domain. This approach ensures the tensor-product structure of the solution space. However, since in this example we do not approximate the solution itself, but only a scalar QoI, this transformation is not required.

The stochastic parametric waveguide model presented here will be used in Sections 3, 4 and 5 to verify the advantages of the considered adaptive methods over their non-adaptive counterparts. Since the model is a low-dimensional one with $N = 6$ parameters, standard UQ approaches can still be applied and compared against adaptive methods. On the contrary, the use of non-adaptive methods will not be possible in Section 6, where moderately high-dimensional models are examined. Whenever the present waveguide model is revisited, the operating frequency is set to 6 GHz. The QoI of interest is chosen to be the absolute value of the reflection coefficient, $|S_{11}|$.

2.2 Uncertainty Propagation and Quantification

As briefly explained in Section 2.1, we consider forward UQ, i.e. uncertainty propagation, which deals with the quantification of the impact of input uncertainties upon the considered QoIs. In contrast, backward or inverse UQ refers to problems where available results regarding the QoI, e.g. measurement values, are used to model, correct, or calibrate input uncertainty. In the following, we focus on uncertainty propagation only.

Following [87, 119], we classify the objectives of uncertainty propagation into four main categories.

1. **Accuracy assessment** refers to the quantification of confidence regarding the predictions of the computational model. The term "validation" is typically used when computational results are tested against measurement data. The term "verification" refers to determining whether the accuracy of a simulation model is sufficient to represent the underlying mathematical model. The same term is also employed for comparisons among different models.

   Since measurements or experimental data are not available in the context of the present work, we focus on verification tasks. A typical verification task will be measuring the accuracy of a surrogate model $\widetilde{g} \approx g$, compared to the original model g. Moreover, estimations regarding statistical measures of a QoI, e.g. its expected value, its variance, or a sensitivity index, will be verified against reference values, if available.

   The accuracy of the surrogate model is measured with an appropriate vector norm, e.g. ℓ^1, ℓ^2, or ℓ^∞. Given a cross-validation set of input-output pairs $Z_{\text{cv}} = \left\{\mathbf{y}_q, g\left(\mathbf{y}_q\right)\right\}_{q=1}^{Q}$, the corresponding errors are given as

$$\epsilon_{\ell^1} = \sum_{q=1}^{Q} \left| g\left(\mathbf{y}_q\right) - \widetilde{g}\left(\mathbf{y}_q\right) \right|, \tag{2.9}$$

$$\epsilon_{\ell^2} = \sum_{q=1}^{Q} \left(g\left(\mathbf{y}_q\right) - \widetilde{g}\left(\mathbf{y}_q\right) \right)^2, \tag{2.10}$$

$$\epsilon_{\ell^\infty} = \max_{q=1,2,\ldots,Q} \left| g\left(\mathbf{y}_q\right) - \widetilde{g}\left(\mathbf{y}_q\right) \right|. \tag{2.11}$$

   The term "cross-validation error" (ϵ_{cv}) will be used for ϵ_{ℓ^∞}.

The accuracy regarding statistical measure estimations is measured using absolute or relative errors

$$\epsilon_{\text{abs}} = \left| \widetilde{\phi} - \phi_{\text{ref}} \right|, \tag{2.12}$$

$$\epsilon_{\text{rel}} = \frac{\left| \widetilde{\phi} - \phi_{\text{ref}} \right|}{|\phi_{\text{ref}}|}, \tag{2.13}$$

where ϕ_{ref} is the reference value of a statistical measure ϕ and $\widetilde{\phi}$ is an estimate, i.e. an approximate value. The reference values will be typically computed with high-order integration schemes. In the case of a MC integration with Q_{MC} samples, the root-mean-square deviation (RMSD)

$$\text{RMSD}[g] = \sqrt{\frac{\mathbb{V}_{\text{MC}}[g]}{Q_{\text{MC}}}}, \tag{2.14}$$

estimates the accuracy of the MC-based expected value $\mathbb{E}_{\text{MC}}[g]$ compared to the "true" expected value, on average [31, 90]. The RMSD is also known as the root-mean-square error (RMSE). Then, the mean relative error given in (2.13) is expected to stagnate at a value close to the normalized root-mean-square deviation (NRMSD)

$$\text{NRMSD}[g] = \frac{\text{RMSD}}{|\mathbb{E}_{\text{MC}}[g]|}, \tag{2.15}$$

which is also referred to as the coefficient of variation (CoV) of the RMSD.

2. **Variance analysis** tries to characterize the model-based predictions regarding the random QoI in terms of robustness. This characterization is typically based on computing statistical moments of the QoI. In many cases, e.g. for distributions close to the normal one, the first two moments, i.e. the expected value $\mathbb{E}[g]$ and the variance $\mathbb{V}[g]$, are sufficient for this task. The first two moments are given by

$$\mathbb{E}[g] = \int_{\Xi} g(\mathbf{y})\, \varrho(\mathbf{y})\, \mathrm{d}\mathbf{y}, \tag{2.16}$$

$$\begin{aligned} \mathbb{V}[g] &= \int_{\Xi} (g(\mathbf{y}) - \mathbb{E}[g])^2\, \varrho(\mathbf{y})\, \mathrm{d}\mathbf{y} \\ &= \mathbb{E}\left[(g - \mathbb{E}[g])^2\right] \\ &= \mathbb{E}\left[g^2\right] - (\mathbb{E}[g])^2. \end{aligned} \tag{2.17}$$

3. **Risk analysis** refers to computing the probabilities of events where the QoI exceeds some critical value. Typical examples are failure or rare-event probabilities. Specialized risk analysis techniques are available, see e.g. [81] for an overview, but are not considered in this work. Instead, we assume that risk analysis can be performed with sampling-based approaches, where a surrogate model $\widetilde{g}$ replaces the original model g to reduce computational costs. For that use, the approximation accuracy of the surrogate model is a critical factor, therefore we focus on careful examinations of the error metrics (2.9), (2.10), and (2.11). We note that estimating failure probabilities with surrogate-based sampling approaches might be problematic in some cases [91].

4. **Uncertainty management** relates the variability of the random QoI to the various sources of uncertainty, with the goal of prioritizing those sources and possibly reducing their number. Sensitivity analyses [136] are commonly employed for that purpose. In this work, we use a variance-based, global sensitivity analysis, known as the Sobol method [146]. The Sobol method decomposes the variance $\mathbb{V}[g]$ into partial variances attributed to the individual parameters or to certain parameter combinations. The related sensitivity metrics, known as Sobol indices, are computed simply as the fractions of the partial variances over the total variance.

   The method starts with a finite-term decomposition of the random output $g(\mathbf{Y})$, such that

$$g(\mathbf{Y}) = g_0 + \sum_{n=1}^{N} g_n(Y_n) + \sum_{n<m}^{N} g_{nm}(Y_n, Y_m) + \cdots + g_{1,2,\ldots,N}(Y_1, Y_2, \ldots, Y_N), \tag{2.18}$$

   where g_0 is a constant function, g_n is a function of the RV Y_n only, g_{nm} is a function of the combination of RVs Y_n and Y_m, and so forth. The terms of (2.18) are proven to be orthogonal, therefore, $g_0 = \mathbb{E}[g]$. Accordingly, the variance $\mathbb{V}[g]$ is decomposed to

$$\begin{aligned} \mathbb{V}[g] &= \sum_{n=1}^{N} \mathbb{V}[g_n] + \sum_{n<m}^{N} \mathbb{V}[g_{nm}] + \cdots + \mathbb{V}\left[g_{1,2,\ldots,N}\right] \\ &= \sum_{n=1}^{N} \mathbb{V}_n + \sum_{n<m}^{N} \mathbb{V}_{nm} + \cdots + \mathbb{V}_{1,2,\ldots,N} \end{aligned} \tag{2.19}$$

   where $\mathbb{V}_n$ is a partial variance attributed only to RV Y_n, $\mathbb{V}_{mn}$ is a partial variance attributed to the combination of RVs Y_n and Y_m, and so forth. "First-

order" or "main-effect" indices measure the influence of the RV Y_n alone, i.e. with all other parameters regarded as constant, and are given as

$$S_n^{\mathrm{FO}} = \frac{\mathbb{V}_n}{\mathbb{V}[g]}. \tag{2.20}$$

We note that

$$\sum_{n=1}^{N} S_n^{\mathrm{FO}} \leq 1. \tag{2.21}$$

"Total-order" or "total-effect" indices measure the influence of the RV Y_n in combination with any number of the remaining parameters Y_m, $m \neq n$, and are given as

$$S_n^{\mathrm{TO}} = \frac{1}{\mathbb{V}[g]} \left(\mathbb{V}_n + \sum_{\substack{m=1 \\ m \neq n}}^{N} \mathbb{V}_{mn} + \cdots + \mathbb{V}_{1,2,\ldots,n\ldots,N} \right), \tag{2.22}$$

where now holds that

$$\sum_{n=1}^{N} S_n^{\mathrm{TO}} \geq 1. \tag{2.23}$$

Typically, one is interested in first and total-order Sobol indices only, however, indices accounting for contributions of specific parameter combinations can be computed in a similar way. The estimation of Sobol indices can be based on sampling approaches, as in [135, 136]. However, the corresponding algorithms are computationally expensive, e.g. given a set of Q_{SA} random inputs, the algorithm from [135] would require $(2N+2)Q_{\mathrm{SA}}$ model evaluations. Hence, computationally inexpensive surrogate models typically replace the original ones for such tasks. Alternatively, Sobol sensitivity information can be directly derived by model approximations with orthogonal terms, see e.g. [23, 149] or Section 4.3.

A wide variety of methods is available to conduct uncertainty propagation studies. Sampling-based methods, such as Monte Carlo (MC) [31, 90], latin hypercube [94], or importance sampling [3], remain the "workhorse" methods in most fields. Methods based on local expansions are also available, such as the perturbation methods considered in [92, 93, 132]. In this work, we focus on approximation methods based on global polynomials, commonly referred to as spectral UQ methods [87, 162]. Spectral methods are the topic of the next section.

2.3 Spectral Methods for Uncertainty Quantification

The fundamental idea behind spectral methods lies in the approximation a PDE solution by a finite series of orthogonal functions, such as orthogonal polynomials or complex exponentials [85]. Spectral UQ methods [87, 162], in particular, aim at approximating the functional dependence of the QoI on the input random parameters. This approximation takes the form of an M-term polynomial series, such that

$$g(\mathbf{y}) \approx \widetilde{g}(\mathbf{y}) = \sum_{j=1}^{M} s_j \Psi_j(\mathbf{y}), \tag{2.24}$$

where $s_j \in \mathbb{C}$ ($s_j \in \mathbb{R}$ if $g : \Xi \to \mathbb{R}$) are series coefficients and $\Psi_j : \Xi \to \mathbb{R}$ are global multivariate polynomials given as

$$\Psi_j(\mathbf{y}) = \prod_{n=1}^{N} \psi_{j_n}(y_n). \tag{2.25}$$

Postponing specific issues to Chapters 3 and 4, it suffices for now to say that the univariate polynomials ψ_{j_n} are chosen in agreement with the univariate PDFs $\varrho_n(y_n)$. A necessary prerequisite for the successful application of spectral UQ methods is that the input-to-output map g is sufficiently smooth with respect to the parameter vector $\mathbf{y}$.

The global index j can be associated with a vector of indices $(j_1, j_2, \ldots, j_N)$ in a unique way, e.g. as in [6, 108]. In particular, we may define an invertible bijective map $\mathscr{V}$, such that

$$\mathscr{V} : \{j_1\}_{j_1=1}^{J_1} \times \{j_2\}_{j_2=1}^{J_2} \times \cdots \times \{j_N\}_{j_N=1}^{J_N} \mapsto \{j\}_{j=1}^{M}, \tag{2.26}$$

where $j_n = 1, 2, \ldots, J_n$ and $M = J_1 J_2 \cdots J_N$. More commonly, a multi-index $\mathbf{j} = (j_1, j_2, \ldots, j_N) \in \mathbb{N}^N$ is employed. Collecting all multi-indices which participate in the approximation in a multi-index set Λ with cardinality $\#\Lambda = M$, the spectral approximation (2.24) can be equivalently written as

$$g(\mathbf{y}) \approx \widetilde{g}(\mathbf{y}) = \sum_{\mathbf{j} \in \Lambda} s_{\mathbf{j}} \Psi_{\mathbf{j}}(\mathbf{y}). \tag{2.27}$$

We denote with $\mathbb{P}_\Lambda$ the polynomial space associated with the multi-index set Λ, such that

$$\mathbb{P}_\Lambda := \operatorname{span}\left\{\Psi_{\mathbf{j}} : \mathbf{j} \in \Lambda\right\}. \tag{2.28}$$

In essence, we are looking for a polynomial approximation $\widetilde{g}$ such that

$$\widetilde{g} = \underset{\pi \in \mathbb{P}_\Lambda}{\arg\min} \|g - \pi\|, \tag{2.29}$$

in some proper norm. The approximation problem can be formally defined as follows.

Definition 2.1 (Spectral approximation problem)**.** *Let $g(\mathbf{y})$ be a function of the multivariate variable $\mathbf{y} = (y_1, y_2, \ldots, y_N) \in \Xi \subseteq \mathbb{R}^N$ and $\mathbb{P}_\Lambda$ a space of multivariate polynomials of $\mathbf{y}$ defined as in* (2.28) *by the multi-index set Λ with cardinality $\#\Lambda = M$. Then, find a polynomial approximation $\widetilde{g} \in \mathbb{P}_\Lambda$ such that $\|\widetilde{g} - g\| \to 0$ as $M \to \infty$, in a proper norm defined on Ξ.*

The approximation $\widetilde{g}$ is often encountered in the literature as the "response surface" or the "surrogate model". Once available, an inexpensive polynomial surrogate model can replace the original model in sampling-based studies. While the low convergence rate of sampling approaches remains, the low cost of the surrogate enables model evaluations on a large set of random realizations of the input parameters. We note that for such a use, the approximation's accuracy is a critical factor.

Sampling is however not always necessary. Statistical information with respect to the QoI is encoded in the expansion terms of (2.24) and certain statistical measures can be provided by directly post-processing the approximation series. More details on that latter approach will be given in Sections 3.5 and 4.3. Finally, we note that, in multiple dimensions, the employed polynomial spaces (2.28) exhibit a tensor-product structure. This structure can be exploited for in the aforementioned post-processing step, as will be discussed in Section 5.3.

2.4 Downward-Closed Multi-Index Sets

Throughout this work we will often request a multi-index set Λ to be downward-closed (also, monotone or lower). Therefore, we provide here a number of definitions related to the downward-closedness property, to be recalled whenever needed in this work.

Definition 2.2 (Forward neighbors of a multi-index set Λ)**.** *Given a multi-index set Λ, we define the set of forward neighbors Λ^+ as*

$$\Lambda^+ := \{\mathbf{j} + \mathbf{e}_n, \forall \mathbf{j} \in \Lambda, \forall n = 1, 2, \ldots, N\}. \tag{2.30}$$

Definition 2.3 (Backward neighbors of a multi-index set Λ)**.** *Given a multi-index set Λ, we define the set of backward neighbors Λ^- as*

$$\Lambda^- := \{\mathbf{j} - \mathbf{e}_n, \forall \mathbf{j} \in \Lambda, \forall n = 1, 2, \ldots, N \ : \ j_n > 0\}. \tag{2.31}$$

Definition 2.4 (Downward-closed multi-index set Λ)**.** *A multi-index set Λ is said to be downward-closed if and only if $\Lambda^- \subset \Lambda$.*

Definition 2.5 (Admissible neighbors of a downward-closed multi-index set Λ)**.** *Given a downward-closed multi-index set Λ, we define the set of admissible neighbors Λ^{adm} as*

$$\Lambda^{\text{adm}} := \left\{ \mathbf{j} \in \Lambda^+ : \mathbf{j} \notin \Lambda, \{\mathbf{j}\}^- \subset \Lambda \right\}. \tag{2.32}$$

In (2.30) and (2.31), $\mathbf{e}_n = (\delta_{mn})_{1 \le m \le N}$ is the n-th unit vector, with δ_{mn} denoting the Kronecker delta. In (2.32), $\{\mathbf{j}\}^-$ denotes the backward neighbors of a multi-index set which contains a single multi-index $\mathbf{j}$, as defined in (2.31), equivalently, the backward neighbors of the multi-index $\mathbf{j}$.

2.5 Adaptivity

As already mentioned in Chapter 1, in this work we employ adaptive approximations which take advantage of parameter anisotropies with respect to their impact upon a QoI. As a result, the discretization of the parameter space will be more refined in certain regions and less refined in others. Adaptive methods of similar nature have been explored since many years in the context of finite element (FE) analysis for anisotropic spatial discretization, see e.g. [1, 158] and the references therein.

Assuming an already available approximation, the main question is how to identify the regions in need of further refinement, as to invest computational resources in an anisotropic, yet meaningful way. A priori error analyses are often employed for that purpose, see e.g. [10, 114] for such approaches in the contexts of stochastic Galerkin and collocation. However, a priori error estimation methods are typically based on strict theoretical assumptions, which rarely hold in most settings. For that reason, the need arises for error estimates which can be computed using already available approximations and then employed to guide adaptivity. Such a posteriori error estimates were first suggested in the FEM context in [7]. Nowadays, a posteriori error estimation has become a field in itself and has been employed for a wide variety of applications, e.g. for quadrature [60] and UQ [29].

The adaptive approximation methods used in this work are based exactly on this idea. In particular, using a readily available approximation of an input-to output map g, we identify possible refinement "candidates", i.e. directions in the parameter space in which the approximation could be further refined, e.g. by adding polynomials of higher order. Each candidate is associated with a local error indicator which measures its contribution to the approximation. The natural choice is to refine the approximation in the direction corresponding to the maximum contribution. A typical example of this procedure is the dimension-adaptive algorithm introduced in [60].

Algorithm 2.1: General adaptive algorithm.

Data: map g, tolerance ϵ, budget B
Result: approximation $\widetilde{g}$

```
Construct initial approximation g̃.
repeat
    Find refinement candidates.
    Compute local error indicators for all candidates.
    Refine g̃ using the candidate with the maximum contribution.
until TERMINATION[ε, B];
```

This procedure is depicted in Algorithm 2.1. The adaptive refinement stops either if a computational budget is reached or if further refinements do not yield any accuracy improvement. Obviously, this is a greedy procedure. As such, its convergence cannot be guaranteed in all cases. However, similar greedy algorithms have been found to perform well in many applications. In this work, similar greedy algorithms are presented in Sections 3.3 and 4.2. The adaptive cross approximation methods discussed in Section 5.2 also fall into the same category.

2.6 Concluding Remarks

In this chapter we have introduced a number of preliminaries which are used throughout this work. We explained the general setting of stochastic parametric problems and models, and we presented a relevant example related to EMF simulations. This model problem will be used in subsequent sections for verification purposes. We further introduced the concept of uncertainty propagation and the main objectives of UQ in this context, providing the necessary definitions and error metrics. Spectral approximations, which constitute the basis of the UQ methods considered in the present work, were discussed next. We then defined the downward-closedness property for multi-index sets, which is often revisited throughout this thesis. Related properties and definitions were also presented. Finally, we introduced a general adaptive refinement approach which constitutes the basis of the methods proposed in subsequent chapters.

3 Dimension-Adaptive Stochastic Collocation

In this chapter we present a dimension-adaptive stochastic collocation algorithm based on nested collocation points. We first recall the basic concepts of the stochastic collocation method, both in the univariate and multivariate case. Next we show how hierarchical collocation schemes can be constructed when nested grids of collocation points are used. We then present the general form of an algorithm for the adaptive construction of hierarchical approximations. The nested collocation points of choice, i.e. Leja and Clenshaw-Curtis (CC), as well as their weighted versions, are discussed next. We proceed with the computation of statistical moments, simply by post-processing the collocation-based approximation. Finally, we compare weighted and unweighted Leja and Clenshaw-Curtis (CC)-based stochastic collocation schemes using the stochastic parametric waveguide model presented in Section 2.1.1.

3.1 Stochastic Collocation

In the context of the stochastic collocation method, the spectral approximation (2.24) is constructed by means of interpolation. We denote the interpolation-based approximation of the input-to-output map g with $\mathscr{I}[g]$, such that $g(\mathbf{y}) \approx \widetilde{g}(\mathbf{y}) = \mathscr{I}[g](\mathbf{y})$. The multivariate approximation $\mathscr{I}[g]$ is formed by combinations of univariate interpolation rules. For a sufficiently smooth map g, the interpolation is typically based on Lagrange polynomials[3].

The method employs a set of collocation points $\left\{\mathbf{y}_j\right\}_{j=1}^{M}$ which belong to the image space Ξ. We denote the set of collocation points with $Z = \left\{\mathbf{y}_j\right\}_{j=1}^{M}$ and its cardinality with $\#Z = M$. There is an $1-1$ relation between the number of collocation points and the number of polynomials employed in the approximation, i.e. each collocation point defines a corresponding Lagrange polynomial. The model is evaluated on all collocation points and the values $g\left(\mathbf{y}_j\right)$ are interpolated, resulting in the approximation $\mathscr{I}[g]$. It holds that $g\left(\mathbf{y}_j\right) = \mathscr{I}[g]\left(\mathbf{y}_j\right)$, $\forall \mathbf{y}_j \in Z$, i.e. the approximation is exact on the collocation points.

Assuming that computing the model evaluations $g\left(\mathbf{y}_j\right)$ outweighs any other operations in the construction of the approximation, the computational cost of the stochastic collocation method depends predominantly on $\#Z$. The accuracy of the approximation in between the collocation points depends on the choice of interpolation nodes. The stochastic collocation problem based on Lagrange interpolation can be formally defined as follows [162, Chapter 7].

[3] According to L. Trefethen [154], E. Waring was the first to present Lagrange interpolation in [159], more than 15 years before Lagrange did.

Definition 3.1 (Stochastic collocation based on Lagrange interpolation)**.** *Let* $Z = \{\mathbf{y}_j\}_{j=1}^{M} \subset \Xi$ *be a set of nodes with cardinality* $\#Z = M$ *and* $\{g(\mathbf{y}_j)\}_{j=1}^{M}$ *the set of the corresponding model evaluations. Let* $\mathbb{P} := \operatorname{span}\{L_j, j = 1, 2, \ldots, M\}$ *be the polynomial space of Lagrange polynomials defined by* Z*, where* $L_i(\mathbf{y}_j) = \delta_{ij}$*,* $i, j = 1, 2, \ldots, M$*. Then, find a polynomial approximation* $\widetilde{g} \in \mathbb{P}$ *such that*

1. $\widetilde{g}(\mathbf{y}_j) = g(\mathbf{y}_j)$*,* $\forall \mathbf{y}_j \in Z$*, and*
2. $\|\widetilde{g} - g\| \to 0$ *as* $M \to \infty$*, in a proper norm defined on* Ξ*.*

3.1.1 Univariate Stochastic Collocation

Univariate interpolation rules are used as building blocks for the stochastic collocation method in multiple dimensions. Therefore, let us first consider the case of a single RV Y. In the stochastic collocation context, any given univariate interpolation rule is typically defined by:

- a non-negative integer $i \in \mathbb{N}_0$ called the interpolation level,
- a monotonically increasing "level-to-nodes" function $m : \mathbb{N}_0 \to \mathbb{N}$ which defines the relation between the interpolation level and the number of interpolation nodes $m_i := m(i)$, where $m_0 = m(0) = 1$, and
- a grid of m_i interpolation nodes, denoted by $Z_i = \{y_{i,j}\}_{j=1}^{m_i}$.

Denoting the univariate interpolation operator with $\mathscr{I}_i$, a univariate map $g(y)$ can be approximated as

$$g(y) \approx \widetilde{g}(y) = \mathscr{I}_i[g](y) = \sum_{j=1}^{m_i} g(y_{i,j})\, l_{i,j}(y), \tag{3.1}$$

where $l_{i,j}$ are univariate nodal Lagrange polynomials of degree $p_i = m_i - 1$, defined on the grid Z_i as

$$l_{i,j}(y) := \begin{cases} \displaystyle\prod_{k=1, k\neq j}^{m_i} \frac{y - y_{i,k}}{y_{i,j} - y_{i,k}}, & i \neq 0, \\ 1, & i = 0. \end{cases} \tag{3.2}$$

We denote the corresponding univariate polynomial space with $\mathbb{P}_i$, such that

$$\mathbb{P}_i = \operatorname{span}\{l_{i,j}, j = 1, 2, \ldots, m_i\}. \tag{3.3}$$

As already mentioned, the quality of the approximation depends crucially on the interpolation nodes, which are chosen according to the univariate PDF. For example, Gauss-Legendre and Gauss-Hermite quadrature nodes are common choices for uniform and normal distributions, respectively.

3.1.2 Tensor-Product Stochastic Collocation

Moving to multiple parameters/dimensions, the simplest form of multivariate collocation consists of tensor product (TP) combinations of univariate interpolation grids $Z_{i_n}^n$ and operators $\mathscr{I}_{i_n}^n$, $n = 1, 2, \ldots, N$. For the ease of presentation and lighter notation, we shall assume that all RVs are identically distributed, such that $Z_{i_n}^n = Z_{i_n}$ and $\mathscr{I}_{i_n}^n = \mathscr{I}_{i_n}$. The extension to the more general case is straightforward.

We introduce the multi-index $\mathbf{i} = (i_1, i_2, \ldots, i_N) \in \mathbb{N}_0^N$ which contains the interpolation levels of all parameters. Generally, the single indices $i_n, n = 1, 2, \ldots, N$, can have different values from one another. The special case where $i_1 = i_2 = \cdots = i_N = i$ is called "isotropic" TP collocation. The multivariate nodes $\mathbf{y_{i,j}} = \left(y_{i_1,j_1}, y_{i_2,j_2}, \ldots, y_{i_N,j_N}\right)$ form the tensor grid

$$\begin{aligned} Z_{\mathbf{i}} &:= \bigtimes_{n=1}^{N} Z_{i_n} \\ &= Z_{i_1} \times Z_{i_2} \times \cdots \times Z_{i_N} \\ &= \left\{y_{i_1,j_1}\right\}_{j_1=1}^{m_{i_1}} \times \left\{y_{i_2,j_2}\right\}_{j_2=1}^{m_{i_2}} \times \cdots \times \left\{y_{i_N,j_N}\right\}_{j_N=1}^{m_{i_N}}, \end{aligned} \tag{3.4}$$

with cardinality $\#Z_{\mathbf{i}} = m_{i_1} m_{i_2} \cdots m_{i_N}$. Denoting the corresponding multivariate Lagrange interpolation operator with $\mathscr{I}_{\mathbf{i}}$, the TP approximation reads

$$\begin{aligned} g(\mathbf{y}) \approx \widetilde{g}(\mathbf{y}) &= \mathscr{I}_{\mathbf{i}}[g](\mathbf{y}) \\ &= \left(\mathscr{I}_{i_1} \otimes \mathscr{I}_{i_2} \otimes \cdots \otimes \mathscr{I}_{i_N}\right)[q](\mathbf{y}) \\ &= \sum_{j_1=1}^{m_{i_1}} \sum_{j_2=1}^{m_{i_2}} \cdots \sum_{j_N=1}^{m_{i_N}} g\left(y_{i_1,j_1}, y_{i_2,j_2}, \ldots, y_{i_N,j_N}\right) \prod_{n=1}^{N} l_{i_n,j_n}(y_n) \\ &= \sum_{\mathbf{y_{i,j}} \in Z_{\mathbf{i}}} g\left(\mathbf{y_{i,j}}\right) L_{\mathbf{i,j}}(\mathbf{y}), \end{aligned} \tag{3.5}$$

where $\otimes$ denotes a tensor product and the multivariate Lagrange polynomials are defined as

$$L_{\mathbf{i,j}}(\mathbf{y}) := \prod_{n=1}^{N} l_{i_n,j_n}(y_n). \tag{3.6}$$

The corresponding TP polynomial space $\mathbb{P}_{\mathbf{i}}$ is given by

$$\mathbb{P}_{\mathbf{i}} = \bigotimes_{n=1}^{N} \mathbb{P}_{i_n}. \tag{3.7}$$

While simple in its conception and construction, the TP stochastic collocation method has a complexity of $\mathcal{O}\left(m_\ell^N\right)$, where $m_\ell = \max_n m_{i_n}$, and therefore becomes intractable even for a relatively small number of parameters. The curse of dimensionality is particularly evident in the isotropic TP case, where the number of collocation points and corresponding model evaluations is exactly equal to m_i^N.

3.1.3 Stochastic Collocation on Smolyak Sparse Grids

Mitigating the complexity to $\mathcal{O}\left(m_\ell \left(\log m_\ell\right)^{N-1}\right)$ while only mildly compromising the approximation's accuracy is possible by employing sparse grids [27]. Sparse grids were first introduced by S. A. Smolyak in [145] for multivariate integration and interpolation purposes and have been further analyzed in the context of the stochastic collocation method in a large number of works, see e.g [6, 13, 108, 116, 163].

We introduce the non-negative integer $\ell \in \mathbb{N}_0$, called the approximation level, and the corresponding Smolyak multi-index set $\Lambda_\ell^{\mathrm{SM}}$, such that

$$\Lambda_\ell^{\mathrm{SM}} \coloneqq \{\mathbf{i} : |\mathbf{i}| \coloneqq i_1 + i_2 + \cdots + i_N \leq \ell\}. \tag{3.8}$$

Further introducing the difference operators

$$\Delta_i \coloneqq \mathscr{I}_i - \mathscr{I}_{i-1}, \tag{3.9}$$

$$\Delta_{\mathbf{i}} \coloneqq \Delta_{i_1} \otimes \Delta_{i_2} \otimes \cdots \otimes \Delta_{i_N}, \tag{3.10}$$

where $\mathscr{I}_{-1}$ is the null operator, the Smolyak approximation formula reads

$$g(\mathbf{y}) \approx \widetilde{g}(\mathbf{y}) = \mathscr{I}_\ell^{\mathrm{SM}}[g](\mathbf{y}) = \sum_{|\mathbf{i}| \leq \ell} \Delta_{\mathbf{i}}[g](\mathbf{y}). \tag{3.11}$$

The Smolyak sparse grid of multivariate interpolation nodes Z_ℓ^{SM} is given as a combination of tensor grids, such that

$$Z_\ell^{\mathrm{SM}} = \bigcup_{\ell - N + 1 \leq |\mathbf{i}| \leq \ell} Z_{\mathbf{i}}. \tag{3.12}$$

In its core, the Smolyak approximation formula (3.11) is nothing more than a linear combination of TP collocation formulas (3.5), where the corresponding tensor grids have a relatively small cardinality.

The original Smolyak sparse grids are of isotropic nature due to (3.8). It is easy to observe that this constraint includes equal contributions from all parameters, that is, equal maximum interpolation levels and number of univariate nodes. As illustrated in Figure 3.1, the complexity reduction from $\mathcal{O}\left(m_\ell^N\right)$ to $\mathcal{O}\left(m_\ell \left(\log m_\ell\right)^{N-1}\right)$

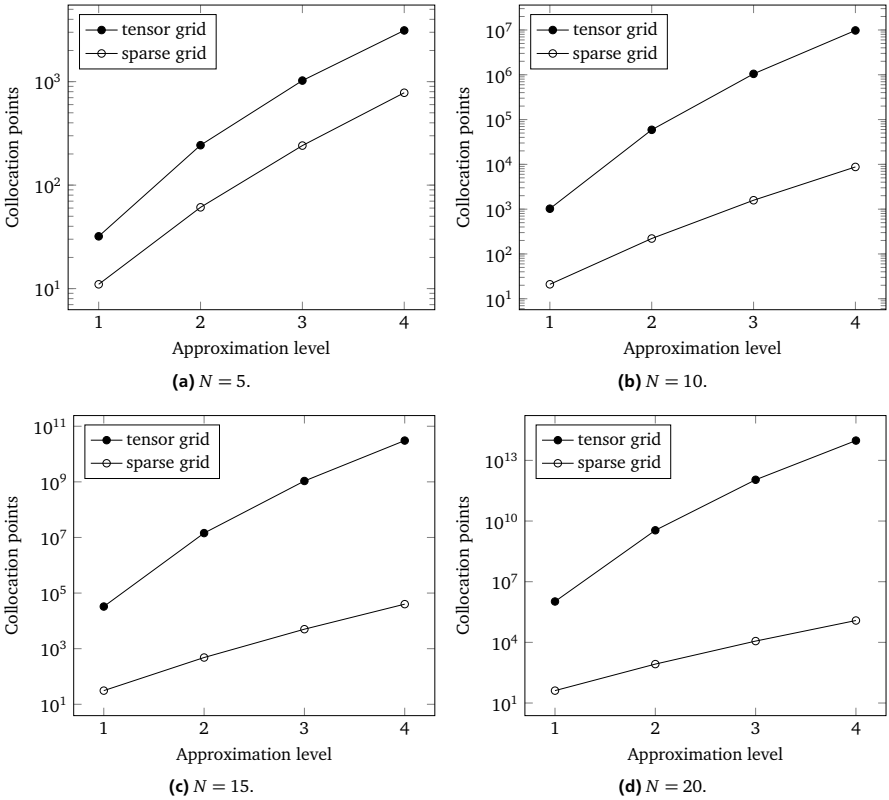

Figure 3.1: Comparison of the complexity growth between tensor grids and sparse grids. In all 4 cases, both grids are isotropic and based on Gauss-Legendre nodes.

due to isotropic sparse grids is tremendous, however, the method is not free of the curse of dimensionality.

In practice, a QoI is rarely equally sensitive to variations from all considered parameters. This anisotropy can be exploited to construct anisotropic sparse grids [117] of reduced sizes, also referred to as "generalized Smolyak" sparse grids [6, 117, 139], as will be shown in Section 3.3.

3.2 Nested Nodes and Hierarchical Interpolation

In this section we explore hierarchical interpolation schemes [37, 83, 139] based on nested grids of interpolation nodes. Employing such schemes, a Smolyak approximation (3.11) can be constructed with reduced costs. Moreover, hierarchical schemes are the basis of the dimension-adaptive stochastic collocation method discussed in Section 3.3.

We start with the observation that Smolyak's formula (3.11) does not necessarily satisfy the interpolation property $g\left(\mathbf{y}_{\mathbf{i},\mathbf{j}}\right) = \mathscr{I}_{\ell}^{\mathrm{SM}}[g]\left(\mathbf{y}_{\mathbf{i},\mathbf{j}}\right), \forall \mathbf{y}_{\mathbf{i},\mathbf{j}} \in Z_{\ell}^{\mathrm{SM}}$. The property is satisfied only if Smolyak's formula is based on nested univariate interpolation nodes, such that $Z_{i-1} \subset Z_i$ [13, Proposition 6]. Common choices are Clenshaw-Curtis (CC)[4] [39, 153] and Gauss-Kronrod-Patterson (GKP) [86, 126] nodes, when the RVs follow uniform and normal distributions, respectively. CC rules for non-uniform weight functions, equivalently, non-uniform PDFs, have been suggested in [148]. Recent works [37, 110] have proposed and investigated the use of Leja points [89] for stochastic collocation purposes. Leja and CC nodes constitute the main choices of nested interpolation nodes in this work and are presented in more detail in Section 3.4.

An illustration of both nested and non-nested univariate interpolation nodes is given in Figure 3.2, where we consider a uniformly distributed parameter in the value range $[-1, 1]$. For both considered cases, i.e. the nested Leja nodes and the non-nested Gauss-Legendre nodes, we consider the level-to-nodes function $m(i) = i + 1$ for $i = 1, 2, 3, 4$. As can be observed from Figure 3.2a, the nodes which correspond to the interpolation levels $i > 0$ include the nodes of all previous levels. This is not true in the non-nested case, presented in Figure 3.2b, where the nodes change for every level refinement.

Irrespective of the choice of the nested nodes, we may employ the corresponding nested grids to derive a hierarchical interpolation scheme. We first define the univariate hierarchical polynomials $h_{i,j}$ as

$$h_{i,j}(y) = \begin{cases} l_{i,j}(y), & y_{i,j} \in Z_i \setminus Z_{i-1}, \\ h_{i-1,j}(y), & y_{i,j} \notin Z_i \setminus Z_{i-1}, \end{cases} \tag{3.13}$$

where $l_{i,j}$ are the Lagrange polynomials defined in (3.2). A comparison between Lagrange and hierarchical polynomials is presented in Figure 3.3. The univariate

[4] Although "Clenshaw-Curtis" is currently the standard term, L. Fejér had already proposed this rule in [54].

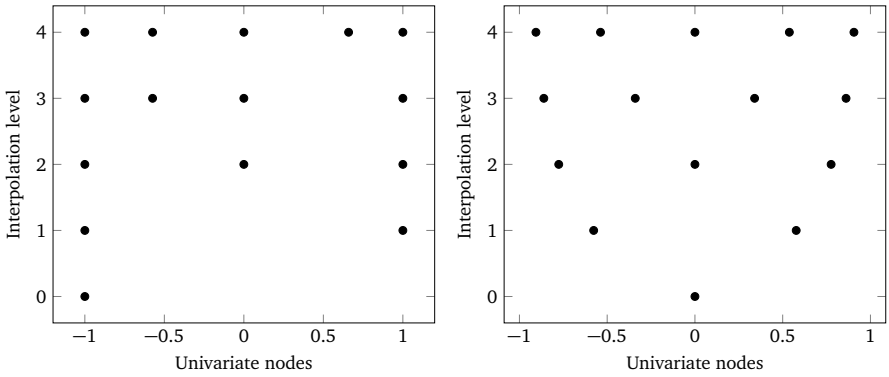

(a) Leja (nested). **(b)** Gauss-Legendre (non-nested).

Figure 3.2: Nested and non-nested univariare interpolation nodes for a uniformly distributed parameter in the range $[-1,1]$ and interpolation levels $i = 1,\dots,4$.

interpolation formula (3.1) can now be equivalently written in the hierarchical format

$$\begin{aligned}\mathscr{I}_i[g](y) &= \mathscr{I}_{i-1}[g](y) + \sum_{y_{i,j}\in Z_i\setminus Z_{i-1}} \Big(g\left(y_{i,j}\right) - \mathscr{I}_{i-1}[g]\left(y_{i,j}\right)\Big)h_{i,j}(y) \qquad (3.14)\\ &= \mathscr{I}_{i-1}[g](y) + \sum_{y_{i,j}\in Z_i\setminus Z_{i-1}} s_{i,j}h_{i,j}(y),\end{aligned}$$

where the coefficients $s_{i,j}$ are known as the "hierarchical surpluses" and are given by

$$s_{i,j} = g\left(y_{i,j}\right) - \mathscr{I}_{i-1}[g]\left(y_{i,j}\right). \qquad (3.15)$$

The hierarchical surpluses $s_{i,j}$ can be interpreted as the contributions of the interpolation nodes $y_{i,j} \in Z_i \setminus Z_{i-1}$ to the already available approximation $\mathscr{I}_{i-1}[g]$. As can be seen from (3.14), an obvious advantage of using nested nodes is that g has to be evaluated only on the new nodes $y_{i,j} \in Z_i \setminus Z_{i-1}$. A further advantage of the hierarchical format (3.14) is that the basis polynomials do not change each time new nodes are added.

In the multivariate case, nested grids of nodes can be constructed by enforcing the use of downward-closed multi-index sets (see Definition 2.4). Downward-closed sets are known to preserve the telescopic properties of the series in (3.11)

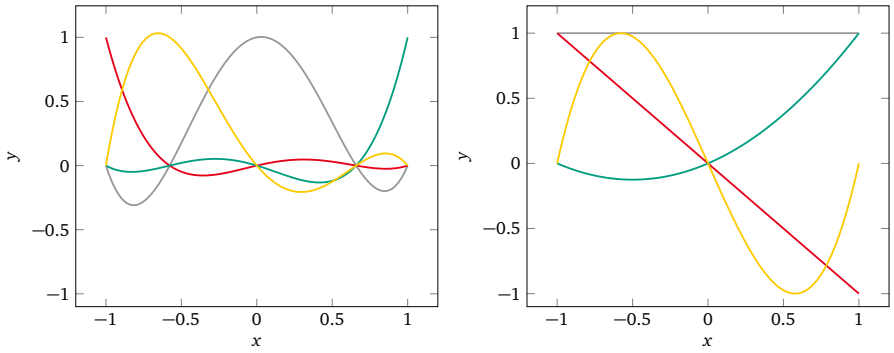

(a) Lagrange polynomials. **(b)** Hierarchical polynomials.

Figure 3.3: Univariate Lagrange and hierarchical polynomials, constructed with the nested sequence of Leja nodes $Z = \{0, -1, 1, -0.577, 0.658\}$.

[60]. Moreover, nested sequences of downward-closed multi-index sets $(\Lambda_k)_{k\geq 0}$, such that $\Lambda_{k-1} \subset \Lambda_k$, $k \geq 1$, result in polynomial approximations of increasing accuracy [37]. Under the condition that the single indices i_n, $n = 1, 2, \ldots, N$, define nested grids of univariate nodes $Z_{i_n - 1} \subset Z_{i_n}$, a sequence of nested downward-closed multi-index sets $(\Lambda_k)_{k\geq 0}$ leads naturally to a sequence of nested grids $\left(Z_{\Lambda_k}\right)_{k\geq 0}$, such that $Z_{\Lambda_{k-1}} \subset Z_{\Lambda_k}$, $\forall k \geq 1$. The corresponding grids of multivariate nodes are given by

$$Z_{\Lambda_k} = \bigcup_{\mathbf{i}\in\Lambda_k} Z_{\mathbf{i}}. \tag{3.16}$$

We note that the Smolyak multi-index set defined in (3.8) also satisfies the downward-closedness property. A sequence of downward-closed multi-index sets $(\Lambda_k)_{k\geq 0}$ can be easily constructed by expanding a readily available multi-index set Λ_k with admissible multi-indices $\mathbf{i} \in \Lambda_k^{\text{adm}}$ (see Definition 2.4).

Applying the Smolyak formula (3.11) in combination with the nested sequence $(\Lambda_k)_{k\geq 0}$ results in the hierarchical multivariate interpolation scheme

$$\begin{aligned} \mathscr{I}_{\Lambda_k}[g](\mathbf{y}) &= \sum_{\mathbf{i}\in\Lambda_k} \Delta_{\mathbf{i}}[g](\mathbf{y}) \\ &= \mathscr{I}_{\Lambda_{k-1}}[g](\mathbf{y}) + \sum_{y_{\mathbf{i},\mathbf{j}}\in Z_{\Lambda_k}\setminus Z_{\Lambda_{k-1}}} \left(g\left(y_{\mathbf{i},\mathbf{j}}\right) - \mathscr{I}_{\Lambda_{k-1}}[g]\left(y_{\mathbf{i},\mathbf{j}}\right)\right) H_{\mathbf{i},\mathbf{j}}(\mathbf{y}) \\ &= \mathscr{I}_{\Lambda_{k-1}}[g](\mathbf{y}) + \sum_{y_{\mathbf{i},\mathbf{j}}\in Z_{\Lambda_k}\setminus Z_{\Lambda_{k-1}}} s_{\mathbf{i},\mathbf{j}} H_{\mathbf{i},\mathbf{j}}(\mathbf{y}), \end{aligned} \tag{3.17}$$

where the multivariate hierarchical polynomials $H_{\mathbf{i},\mathbf{j}}$ are given as

$$H_{\mathbf{i},\mathbf{j}}(\mathbf{y}) = \prod_{n=1}^{N} h_{i_n,j_n}(y_n), \tag{3.18}$$

and the hierarchical surpluses $s_{\mathbf{i},\mathbf{j}}$ are given, similarly to (3.15), by

$$s_{\mathbf{i},\mathbf{j}} = g\left(y_{\mathbf{i},\mathbf{j}}\right) - \mathscr{I}_{\Lambda_{k-1}}[g]\left(y_{\mathbf{i},\mathbf{j}}\right). \tag{3.19}$$

As in the univariate case, the main advantage of the hierarchical scheme (3.17) is the re-use of the already computed model evaluations for the sets $\Lambda_0, \Lambda_1, \ldots, \Lambda_{k-1}$.

The remaining question is how to construct sequences of nested downward-closed multi-index sets $(\Lambda_k)_{k\geq 0}$, that is, which multi-indices should be added to a currently available multi-index set Λ_k. This will be the topic of the next section.

3.3 Dimension-Adaptive Stochastic Collocation

Our goal is to construct nested sequences of downward-closed sets $(\Lambda_k)_{k\geq 0}$, such that the resulting anisotropic (generalized Smolyak) approximations capture the parameter anisotropy with respect to the input-to-output map g. To that end, we employ a dimension-adaptive algorithm which originates from the seminal work of Gerstner and Griebel in [60] for quadrature purposes, and has been subsequently used in a large number of works for both quadrature and interpolation, see e.g. [37, 83, 110, 139, 140].

Let us assume that a multivariate hierarchical approximation (3.17) based on the downward-closed multi-index set Λ_k, from now on called the "activated" set, is readily available. Considering the admissible set Λ_k^{adm} defined as in (2.32), it can be easily verified that expanding Λ_k with a multi-index $\mathbf{i} \in \Lambda_k^{\mathrm{adm}}$ ensures that the downward-closedness property holds (see Definition 2.4). Each multi-index $\mathbf{i} \in \Lambda_k^{\mathrm{adm}}$ is associated with a set of collocation points $\mathbf{y}_{\mathbf{i},\mathbf{j}} \in Z_{\mathbf{i}} \setminus Z_{\Lambda_k}$. In turn, hierarchical surpluses $s_{\mathbf{i},\mathbf{j}}$, defined as in (3.19), can be associated to those collocation points. Motivated by the interpretation of hierarchical surpluses as contributions to the already available approximation, we define the contribution indicators $\eta_{\mathbf{i}}$, $\mathbf{i} \in \Lambda_k^{\mathrm{adm}}$ as

$$\eta_{\mathbf{i}} := \frac{1}{\#\left(Z_{\mathbf{i}} \setminus Z_{\Lambda_k}\right)} \sum_{\mathbf{y}_{\mathbf{i},\mathbf{j}} \in Z_{\mathbf{i}} \setminus Z_{\Lambda_k}} |s_{\mathbf{i},\mathbf{j}}|, \tag{3.20}$$

where the total contribution of a multi-index is scaled by the respective costs, i.e. model evaluations on the extra nodes. The activated set is then expanded by adding the admissible multi-index which corresponds to the maximum contribution $\eta_{\mathbf{i}}$, i.e.

$$\Lambda_{k+1} = \Lambda_k \cup \{\mathbf{i}_{k+1}\}, \quad \text{where} \quad \mathbf{i}_{k+1} = \operatorname*{arg\,max}_{\mathbf{i} \in \Lambda_k^{\mathrm{adm}}} \eta_{\mathbf{i}}. \tag{3.21}$$

This procedure is continued iteratively until some stopping criteria are met, e.g. a certain simulation budget B is reached or the total or maximum contribution of the admissible set drops below a pre-defined tolerance ϵ. The stopping criteria are formally given as

$$\#Z_{\Lambda_k} + \#Z_{\Lambda_k^{\text{adm}}} \geq B, \tag{3.22}$$

$$\sum_{\mathbf{i} \in \Lambda_k^{\text{adm}}} \eta_{\mathbf{i}} \leq \epsilon, \quad \text{or} \tag{3.23a}$$

$$\max_{\mathbf{i} \in \Lambda_k^{\text{adm}}} \eta_{\mathbf{i}} \leq \epsilon. \tag{3.23b}$$

Once a stopping criterion is met, the final approximation is constructed using all multi-indices $\mathbf{i} \in \Lambda_k \cup \Lambda_k^{\text{adm}}$, in order to include the already computed contributions of the admissible multi-indices. The starting multi-index set Λ_0 is typically chosen to be $\Lambda_0 = \{(0, 0, \ldots, 0)\}$. However, the procedure can be initialized with any downward-closed set and can therefore be restarted at will.

The dimension-adaptive procedure explained above is presented in Algorithm 3.1. It is a greedy algorithm, the costs of which at any given step k are equal to the number of collocation points $\#Z_\Lambda = \#Z_{\Lambda_k} + \#Z_{\Lambda_k^{\text{adm}}}$, assuming that the dominating cost factor is the number of model evaluations $g\left(\mathbf{y}_{\mathbf{i},\mathbf{j}}\right)$, $\mathbf{y}_{\mathbf{i},\mathbf{j}} \in Z_\Lambda$. In the worst case, i.e. if all parameters affect the QoI in a similar way, the greedy algorithm is expected to construct the isotropic Smolyak multi-index set Λ_ℓ^{SM}, defined in (3.8).

3.4 Nested Collocation Points

As already pointed out, adaptive sparse grids are based on nested univariate collocation grids. Therefore, nestedness is a key requirement for the employed collocation points. Moreover, the Lebesgue constant associated with the collocation points must remain bounded such that the interpolation yields accurate results [37]. Finally, the selected points must form accurate quadrature rules to be used for the computation of statistical measures, as will be discussed in 3.5. In this work we focus on two families of collocation points which satisfy all three requirements, namely the Clenshaw-Curtis and Leja nodes, described in the following.

3.4.1 Leja

The first option is to base the collocation on Leja sequences [89], which have recently been utilized in the context of sparse grids for interpolation and integration

Algorithm 3.1: Dimension-adaptive stochastic collocation.

Data: map $g(\mathbf{y})$, initial activated multi-index set Λ_0, tolerance ϵ, budget B
Result: sparse grid Z_Λ, approximation $\mathscr{I}_\Lambda[g]$

$k = 0$
while TRUE **do**
  Compute admissible set $\Lambda_k^{\text{adm}} := \left\{\mathbf{i} \in \Lambda_k^+ : \mathbf{i} \notin \Lambda_k \quad \text{and} \quad \{\mathbf{i}\}^- \subset \Lambda_k\right\}$.
  Compute hierarchical surpluses $s_{\mathbf{i},\mathbf{j}} := g\left(\mathbf{y}_{\mathbf{i},\mathbf{j}}\right) - \mathscr{I}_{\Lambda_k}[g]\left(\mathbf{y}_{\mathbf{i},\mathbf{j}}\right), \forall \mathbf{y}_{\mathbf{i},\mathbf{j}} \in Z_{\Lambda_k^{\text{adm}}} \setminus Z_{\Lambda_k}$.
  Compute contribution indicators $\eta_{\mathbf{i}} := \frac{1}{\#\left(Z_{\mathbf{i}} \setminus Z_{\Lambda_k}\right)} \sum_{\mathbf{y}_{\mathbf{i},\mathbf{j}} \in Z_{\mathbf{i}} \setminus Z_{\Lambda_k}} |s_{\mathbf{i},\mathbf{j}}|, \forall \mathbf{i} \in \Lambda_k^{\text{adm}}$.
  Compute current budget $B_k = \#Z_{\Lambda_k} + \#Z_{\Lambda_k^{\text{adm}}}$.
  Compute current total (or maximum) contribution $\epsilon_k = \sum_{\mathbf{i} \in \Lambda_k^{\text{adm}}} \eta_{\mathbf{i}}$ $\left(\text{or } \epsilon_k = \max_{\mathbf{i} \in \Lambda_k^{\text{adm}}} \eta_{\mathbf{i}}\right)$.
  if $B_k \geq B$ OR $\epsilon_k \leq \epsilon$ **then**
    break loop
  end
  Find new multi-index $\mathbf{i}_{k+1} = \arg\max_{\mathbf{i} \in \Lambda_k^{\text{adm}}} \eta_{\mathbf{i}}$.
  Update activated set $\Lambda_{k+1} = \Lambda_k \cup \{\mathbf{i}_{k+1}\}$.
  $k = k + 1$
end
Construct final multi-index set $\Lambda = \Lambda_k \cup \Lambda_k^{\text{adm}}$, sparse grid Z_Λ and approximation $\mathscr{I}_\Lambda[g]$.

purposes in [37, 67, 110, 114]. A sequence of standard, unweighted Leja nodes $(y_k)_{k \geq 0}$, $y_k \in [-1, 1]$, is defined recursively by solving for each node y_k, $k \geq 1$, the optimization problem

$$y_{k+1} = \underset{y \in [-1,1]}{\arg\max} \prod_{j=0}^{k} \left|y - y_j\right|. \tag{3.24}$$

The initial node $y_0 \in [-1, 1]$ is chosen arbitrarily. In this work, we focus on the so-called weighted Leja nodes [110], the definition of which involves a continuous and positive weight function, here given by a univariate PDF $\varrho(y)$, $\varrho : \Xi \to \mathbb{R}_+$. A sequence of univariate weighted Leja nodes $(y_k)_{k \geq 0}$, $y_k \in \Xi$, can be constructed by solving the optimization problem

$$y_{k+1} = \underset{y \in \Xi}{\arg\max} \sqrt{\varrho(y)} \prod_{j=0}^{k} \left|y - y_j\right|, \tag{3.25}$$

where the starting node $y_0 \in \Xi$ is again arbitrarily chosen. When the weight function $\varrho(y)$ coincides with the PDF of the uniform distribution $\mathscr{U}[-1,1]$, problem (3.25) reduces to the problem (3.24) of the unweighted case. We note that weighted Leja sequences can be constructed also on unbounded domains, i.e. when normal (Gaussian) weight functions are considered. The properties of weighted Leja nodes on unbounded domains have been investigated in [80, 110]. In [110], an analysis with respect to other weight functions is also available, e.g. beta or gamma.

The choice of Leja nodes offers a number of advantages. First of all, the optimization problems (3.24) and (3.25) result in sequences of nodes which are nested by construction. Secondly, Leja nodes allow complete freedom in the choice of the level-to-nodes function $m(i)$, e.g. we may consider $m(i) = i+1$, $i \in \mathbb{N}_0$, in order to get the minimum of one extra node per interpolation level refinement. In comparison, CC or GKP nodes would restrict us to the rapidly growing level-to-nodes functions $m(i) = 2^i + 1$ and $m(i) = 2^i - 1$, respectively. Third, by integrating an interpolant constructed with a weighted Leja sequence (see Section 3.5), a Leja-based quadrature rule can be derived [110]. Finally, due to their definition in (3.25), weighted Leja nodes are tailored to the given PDF.

3.4.2 Clenshaw-Curtis

The second option is to use as univariate collocation points the nodes of the Clenshaw-Curtis (CC) quadrature rule. The rule has been proposed in [39] for the integral approximation

$$\int_{-1}^{1} g(y)\mathrm{d}y \approx \sum_{j=1}^{m_i} w_{i,j}\, g\left(y_{i,j}\right), \tag{3.26}$$

where $y_{i,j}$ are the quadrature nodes and $w_{i,j}$ the quadrature weights of a level-i univariate quadrature rule. The standard CC nodes are the extrema of the Chebyshev polynomials $T_k(y)$ in the interval $[-1,1]$, plus the boundary points of the interval [153]. The weights are typically computed by sums of trigonometric functions [148]. Nested CC nodes are obtained with the level-to-nodes function $m_i = m(i) = 2^i + 1$, with $m_0 = m(0) = 1$, such that $Z_{i-1} \subset Z_i$, with $\#\left(Z_i \setminus Z_{i-1}\right) = 2^i$.

For integrations over general bounded domains $[a,b]$, the quadrature nodes and weights can be easily derived by simply scaling the nodes and weights in $[-1,1]$, such that

$$y'_{i,j} = \frac{b-a}{2} y_{i,j} + \frac{a+b}{2}, \qquad w'_{i,j} = \frac{b-a}{2} w_{i,j}, \tag{3.27}$$

where $y'_{i,j}, w'_{i,j}$ refer to the interval $[a,b]$. It is therefore straightforward to extend the quadrature to integrals $\int_a^b q(y)\varrho(y)\mathrm{d}y$ with a constant weight function, e.g. a uniform PDF $\varrho(y) = 1/(b-a)$ with support in $[a,b]$, such that

$$\begin{aligned}\int_a^b g(y)\varrho(y)\mathrm{d}y &= \int_a^b g(y)\frac{1}{b-a}\mathrm{d}y \\ &\approx \frac{1}{b-a}\sum_{j=1}^{m_i} w'_{i,j}\, g\left(y'_{i,j}\right) \\ &= \frac{1}{b-a}\sum_{j=1}^{m_i}\frac{b-a}{2} w_{i,j}\, g\left(y'_{i,j}\right) \\ &= \frac{1}{2}\sum_{j=1}^{m_i} w_{i,j}\, g\left(\frac{b-a}{2}y_{i,j} + \frac{a+b}{2}\right). \end{aligned} \tag{3.28}$$

In the case of a non-uniform PDF $\varrho(y)$, or, generally, a non-constant weight function, the quadrature weights must be recomputed. As already said, the nodes correspond to extrema of Chebyshev polynomials, and are therefore independent of the weight function. A numerically efficient construction of non-uniform CC weights has been given in [148]. The proposed approach is based on the discrete sine/cosine transform and is adopted in this work. To be precise, the j-th CC weight is given by

$$w_{i,j} = \frac{1}{2(j-1)}\left(2\sum_{k=0}^{j-1}(-1)^k\gamma_k + \gamma_0 + (-1)^j\gamma_{j+1}\right), \tag{3.29}$$

where $\gamma_k = \int_{-1}^{1} T_k(y)\varrho(y)\,\mathrm{d}y$ represent moments of the Chebyshev polynomial T_k, to be precomputed.

3.5 Post-processing the Approximation

As already mentioned in Section 2.3, the stochastic collocation approximation can be used as an inexpensive surrogate model and replace the original model in sampling-based computations of statistical measures, e.g. failure probabilities or response PDFs. Another option is to transform the polynomial basis into an orthogonal one [30] and exploit the orthogonality property for the direct derivation of statistical information, as will be explained in Section 4.3. This latter approach is not considered here.

Sampling is not necessary for the computation of statistical moments of the QoI, e.g. its expected value $\mathbb{E}[g]$, variance $\mathbb{V}[g]$, defined in (2.16) and (2.17), respectively, or higher order moments. Moment estimations can be efficiently tackled

with appropriate quadrature schemes, derived after the collocation-based approximation.

First, considering a univariate approximation of level i as in (3.14), thus employing m_i collocation points, we apply the expectation operator such that the expected (mean) value of the QoI can be estimated as

$$\begin{aligned}\mathbb{E}[g] = \int_{\Xi} g(y)\varrho(y)\mathrm{d}y \approx \int_{\Xi} \mathscr{I}_i[g](y)\varrho(y)\mathrm{d}y = \int_{\Xi}\left(\sum_{j=1}^{m_i} s_{i,j}h_{i,j}(y)\right)\varrho(y)\mathrm{d}y \\ = \sum_{j=1}^{m_i} s_{i,j}\int_{\Xi} h_{i,j}(y)\varrho(y)\mathrm{d}y = \sum_{j=1}^{m_i} s_{i,j}\mathbb{E}\left[h_{i,j}\right].\end{aligned} \tag{3.30}$$

Similar schemes can be used for the estimation of higher order moments, using approximations of the quantities g^p, where p denotes the moment order [139]. We observe that (3.30) is similar to a univariate quadrature rule, here denoted with $\mathscr{Q}_i$, where the function's evaluations on the quadrature nodes are incorporated in the coefficients $s_{i,j}$ and the quadrature weights $w_{i,j}$ are given by

$$w_{i,j} = \int_{\Xi} h_{i,j}(y)\,\varrho(y)\,\mathrm{d}y = \mathbb{E}\left[h_{i,j}\right]. \tag{3.31}$$

Moving to the multivariate case, tensor-product quadrature rules can be constructed similarly to (3.5), such that

$$\begin{aligned}\mathscr{Q}_{\mathbf{i}}[g] &= \left(\mathscr{Q}_{i_1}\otimes\mathscr{Q}_{i_2}\otimes\cdots\otimes\mathscr{Q}_{i_N}\right)[g] \\ &= \sum_{j_1=1}^{m_{i_1}}\sum_{j_2=1}^{m_{i_2}}\cdots\sum_{j_N=1}^{m_{i_N}} g\left(y_{i_1,j_1}, y_{i_2,j_2}, \ldots, y_{i_N,j_N}\right)\prod_{n=1}^{N} w_{i_n,j_n} \\ &= \sum_{\mathbf{y}_{\mathbf{i},\mathbf{j}}\in Z_{\mathbf{i}}} g\left(\mathbf{y}_{\mathbf{i},\mathbf{j}}\right) w_{\mathbf{i},\mathbf{j}},\end{aligned} \tag{3.32}$$

where the multivariate weights $w_{\mathbf{i},\mathbf{j}}$ are given as products of the univariate ones, i.e.

$$w_{\mathbf{i},\mathbf{j}} = \prod_{n=1}^{N} w_{i_n,j_n}. \tag{3.33}$$

Then, assuming a readily available multivariate approximation based on a multi-index set with cardinality $\#\Lambda = K$, given as

$$g(\mathbf{y}) \approx \sum_{\mathbf{i}\in\Lambda} \Delta_{\mathbf{i}}[g](\mathbf{y}) = \sum_{k=1}^{K} s_k H_k(\mathbf{y}), \tag{3.34}$$

as well as an 1-1 relation between the global index k and all combinations of the multi-indices $\mathbf{i},\mathbf{j}$, corresponding to the collocation points $\mathbf{y}_{\mathbf{i},\mathbf{j}} \in Z_\Lambda$, the expected value of the QoI can be estimated as

$$\mathbb{E}[g] \approx \int_\Xi \left(\sum_{k=1}^{K} s_k H_k(\mathbf{y}) \right) \varrho(\mathbf{y}) \mathrm{d}\mathbf{y} = \sum_{k=1}^{K} s_k \int_\Xi H_k(\mathbf{y}) \varrho(\mathbf{y}) \mathrm{d}\mathbf{y} = \sum_{k=1}^{K} s_k \mathbb{E}[H_k], \quad (3.35)$$

where the multivariate hierarchical polynomials H_k are given as products of univariate ones, as in (3.18). Taking also into consideration (3.31) and (3.33), it holds that

$$\mathbb{E}[H_k] = \mathbb{E}\left[H_{\mathbf{i},\mathbf{j}}\right] = \prod_{n=1}^{N} \mathbb{E}\left[h_{i_n,j_n}\right] = \prod_{n=1}^{N} w_{i_n,j_n} = w_{\mathbf{i},\mathbf{j}}. \quad (3.36)$$

Therefore, (3.35) is similar to a K-term multivariate quadrature rule, where the function's evaluations on the quadrature nodes are incorporated in the coefficients s_k and the quadrature weights are given as products of univariate weights.

3.6 Application to the Model Problem

We use the stochastic parametric waveguide model presented in Section 2.1.1 to perform a series of numerical experiments for verification purposes. In particular, we compare the stochastic collocation method based on isotropic and adaptive-anisotropic sparse grids, using either CC or Leja nested rules. Two choices regarding the distributions of the input RVs are considered, in particular uniform and skewed beta ones.

In the first study, presented in Section 3.6.1, we verify the advantages of adaptivity as well as the suitability of both node families for a standard choice regarding the PDFs of the input RVs. In the literature, the most commonly used distributions for modeling the input RVs are the uniform and the normal ones. We denote the uniform distribution with $\mathcal{U}(a,b)$, where a and b are the distribution's lower and upper bounds. The normal distribution is denoted with $\mathcal{N}(\mu,\sigma^2)$, where μ is the mean value, σ the standard deviation and σ^2 the variance. The stochastic collocation method has been extensively studied for those distributions and suitable nested and non-nested interpolation nodes, besides the ones considered here, are available. Here, we only consider the uniform case, where studies regarding the performance of CC and Leja-based collocation are already available in the literature [110, 114]. Hence, this setting is suitable for a sanity check, before considering more complicated cases.

In the second study, we consider non-standard choices of input PDFs, i.e. neither uniform nor normal. On the one hand, the choice of uniform PDFs is often

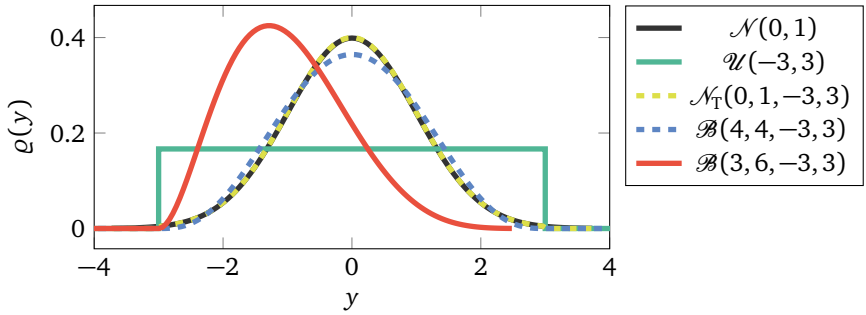

Figure 3.4: Approximations of the standard normal PDF by truncated normal and beta distributions, restricted to the $\pm 3\sigma$ range. The uniform distribution and a non-symmetric, right-skewed beta distribution in the same range are also presented.

criticized as unphysical, in the sense that a stochastic quantity is more naturally described in a Gaussian-like way by a dominant value and a tendency of variation around the dominant value. On the other hand, due to their support on unbounded domains, the use of normal distributions admits the occurrence of unphysical RV realizations, e.g. negative geometry values, with a non-zero probability. To avoid unphysical parameter realizations, certain value boundaries must be respected, i.e. $y \in [a, b]$. Such a modeling can be accomplished by using either a truncated normal distribution $\mathcal{N}_{\mathrm{T}}(\mu, \sigma^2, a, b)$ [28], or a beta distribution $\mathcal{B}(\alpha, \beta, a, b)$ with specific shape parameters α and β [165]. See Figure 3.4 for an illustration. Beta-based approximations of the normal distributions have been considered in our work [57]. We will here consider even more "exotic" distributions, in particular the non-symmetric beta distribution presented in Figure 3.4, and compare the performance of weighted Leja [110] and weighted CC [148] collocation, as presented in Sections 3.4.1 and 3.4.2. Besides our own contribution [96], we are not aware of other studies in this context.

A few words regarding the implementation of the employed stochastic collocation methods:

- The Sparse Grids MATLAB Kit [9, 150] is used for all isotropic Smolyak methods. The same software is also used for the dimension-adaptive scheme based on CC nodes. The software does not support distributions other than normal or uniform, and is therefore complemented by self-developed implementations when beta-distributed input RVs are considered.

- The in-house developed Dimension Adaptive Leja Interpolation (DALI) Python software [95] is used for the dimension-adaptive scheme based on weighted Leja nodes. Our implementation is partially based on the Chaospy Python package [53].

3.6.1 Leja versus Clenshaw-Curtis for Uniform Input Distributions

We assume that all waveguide parameters follow uniform distributions with variations ±10% around their nominal values (see Table 2.3), such that $Y_n \sim \mathcal{U}(a_n, b_n)$, where $a_n = y_n^{\text{nom}} - 0.1 \cdot y_n^{\text{nom}}$ and $b_n = y_n^{\text{nom}} + 0.1 \cdot y_n^{\text{nom}}$, $n = 1, 2, \ldots, 6$. We apply four variants of the stochastic collocation method, namely isotropic Smolyak and dimension-adaptive approximations based on CC and Leja nodes. In the isotropic Smolyak case we employ the approximation levels $\ell_{\text{CC}} = 1, 2, \ldots, 5$, and $\ell_{\text{Leja}} = 1, \ldots, 9$. The relation between approximation levels and costs, i.e. number of collocation points, is presented in Figure 3.5. We remind the reader of the respective level-to-nodes functions, $m_{\text{CC}}(i) = 2^i + 1$ and $m_{\text{Leja}}(i) = i + 1$ (see Section 3.4), which explains the difference in the number of collocation nodes per approximation level. In the dimension-adaptive case we use the simulation budget-based stop criterion (3.22) with an increasing budget and obtain a total of 59 hierarchical approximations (3.17), 9 in the budget range [10,90] and 50 in the budget range [100, 5000].

We first compare all methods with respect to their approximation accuracy for similar costs. The approximation accuracy is measured with the cross-validation error defined in (2.11), on a cross-validation set of size $Q_{\text{cv}} = 10^5$. The results are presented in Figure 3.6. As can be observed, the dimension-adaptive approximations outperform their isotropic counterparts by several orders of magnitude. Moreover, Leja nodes seem to be superior to CC nodes for approximation purposes. This finding is in agreement with the results of the available literature [110, 114].

Next, we measure the accuracy of statistical moment estimations using the relative error metric defined in (2.13). The reference values have been computed up to machine accuracy with all collocation methods. The corresponding results are presented in Figure 3.7. For both the expected value and the variance, the dimension-adaptive estimates converge significantly faster than the isotropic ones, as shown in Figures 3.7a and 3.7b. On the contrary the CC nodes are superior to Leja nodes for the quadrature-based moment estimations, as shown in Figure 3.6,. Again, this result coincides with the observations of the available literature [110, 114].

Finally, we investigate the reasons behind the success of the adaptive methods by examining the sensitivity of the QoI with respect to the input parameters, and the interpolation levels employed by the adaptive collocation methods. Sobol indices [146] are used as sensitivity metrics. The Sobol indices are computed with

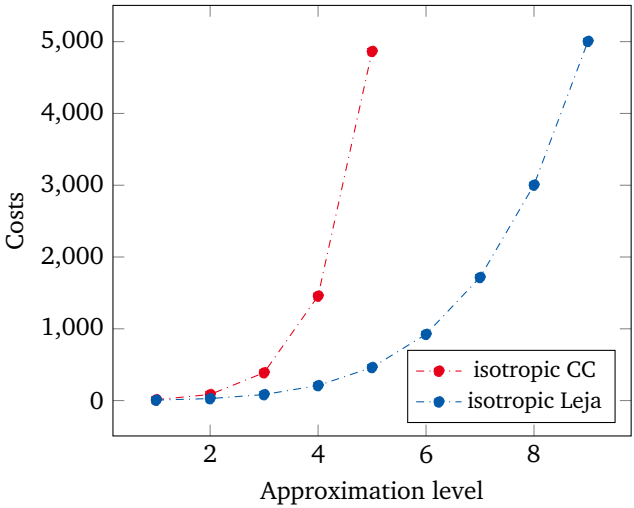

Figure 3.5: Approximation level versus costs for six-dimensional isotropic Smolyak sparse grids based on CC and Leja nodes.

a sampling-based approach which employs Saltelli's algorithm [135], requiring a total of $(2N+2)Q_{\mathrm{SA}} = 14 \cdot 10^4$ model evaluations, where $Q_{\mathrm{SA}} = 10^4$ is the size of the random input realizations set. Both the original model and the most accurate surrogate models are used and their results are compared in order to verify that a sufficiently accurate surrogate model can reliably replace the original model for such tasks. The respective results are presented in Figure 3.8.

Figure 3.8a presents the maximum interpolation level per parameter for the most accurate CC and Leja-based adaptive approximations, i.e. the ones employing (approximately) 5000 polynomials. The difference in the maximum interpolation levels is attributed to the different level-to-nodes functions employed by the CC and Leja-based schemes, i.e. $m_{\mathrm{CC}}(i) = 2^i + 1$ and $m_{\mathrm{Leja}}(i) = i + 1$. The adaptive schemes disregard completely two of the input parameters, i.e. the height h and the vacuum offset d, and refine the grids in the remaining directions. The anisotropy in both adaptive schemes is the same, i.e. the same two parameters are disregarded, the approximation level refinements for parameters l, ε_{r} and μ_{r} are the same, and parameter w is refined more than the rest.

Figure 3.8b presents the Sobol sensitivity index per parameter. Irrespective of the model used, i.e. original, CC-based surrogate, or Leja-based surrogate, the computed Sobol indices are almost identical. Hence, we may say that a sufficiently

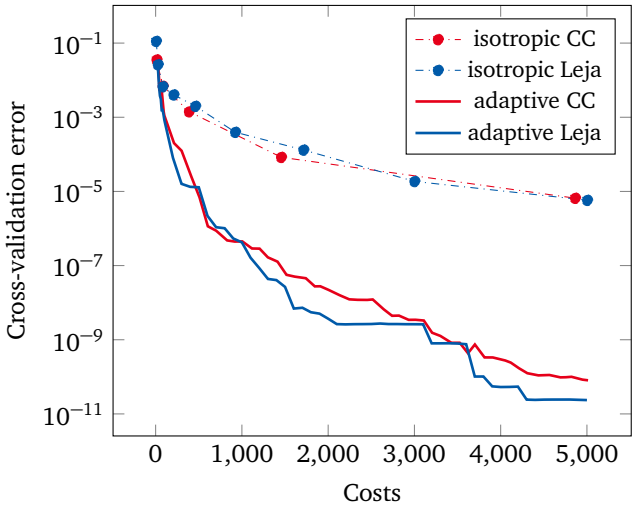

Figure 3.6: Cross-validation error versus costs for the dimension-adaptive stochastic collocation method, applied to the model waveguide problem with uniform random inputs.

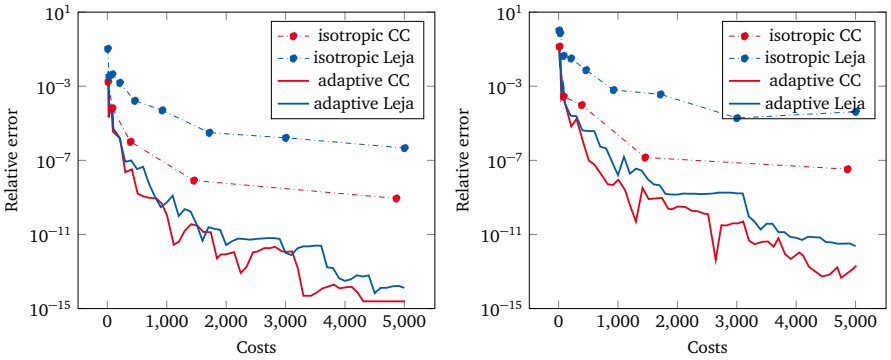

(a) Relative error versus costs for $\mathbb{E}[|S_{11}|]$.

(b) Relative error versus costs for $\mathbb{V}[|S_{11}|]$.

Figure 3.7: Moment relative errors versus costs for the dimension-adaptive stochastic collocation method, applied to the model waveguide problem with uniform random inputs.

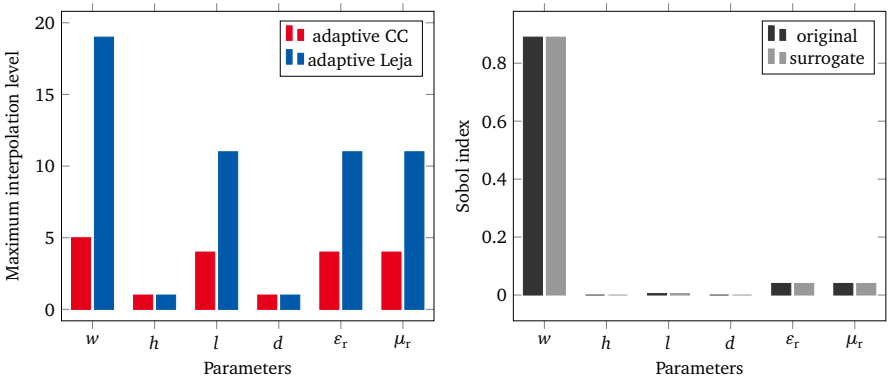

(a) Maximum interpolation level per parameter.

(b) Sobol sensitivity index per parameter.

Figure 3.8: Sensitivity analysis and adaptivity results of the dimension-adaptive stochastic collocation method, applied to the model waveguide problem with uniform random inputs.

accurate surrogate model can reliably substitute the original one in sensitivity analyses. The sum of all first-order Sobol indices, defined in (2.20), is approximately equal to one, therefore, the difference between first-order and total-order Sobol indices, the latter defined in (2.22), is negligible. Hence, in Figure 3.8b we only present the first-order Sobol indices computed with the original model and with the Leja-based approximation. According to the results, almost 90% of the variance of the QoI can be attributed to the waveguide's width w. The two material parameters ε_r and μ_r account for 4% of the output variance each, while the filling length l has a small contribution of approximately 0.5%. The QoI is found to be completely insensitive to variations of the height h and vacuum offset d.

We note here that, with the exception of completely insignificant parameters, maximum interpolation levels do not provide any sensitivity information. For example, we may expect a low interpolation level for a significant parameter which varies in a linear fashion with respect to the QoI, and a higher interpolation level for a less significant parameter with an oscillatory relation. Considering the model waveguide problem and the results presented in Figure 3.8, this can be verified by the maximum interpolation level of parameter l, which is similar to the ones of the material parameters, despite its significantly lower Sobol index value.

3.6.2 Leja versus Clenshaw-Curtis for Non-Symmetric Beta Input Distributions

We now consider that all waveguide parameters follow beta distributions with the shape parameters $\alpha = 3$ and $\beta = 6$ and with support similar to the uniform distributions considered in Section 3.6.1. For the selected shape parameters, the beta distributions are non-symmetric and right-skewed, as shown in Figure 3.4. Leja points and weights tailored to this distribution can be constructed by the definition (3.25). The CC nodes are not affected by the input densities as long as the support remains unchanged, however, the quadrature weights must be modified according to (3.29).

We again consider comparisons between Leja and CC nodes with respect to both interpolation and quadrature accuracy. Since the benefits of adaptivity have already been established in Section 3.6.1, we only use the dimension adaptive algorithm 3.1. Reference moment values are computed up to machine accuracy with both dimension-adaptive schemes. Cross-validation errors are computed with the cross-validation sample of size $Q_{\text{cv}} = 10^5$ employed in Section 3.6.1, i.e. the random realizations are drawn from a uniform-based joint PDF. The selected joint PDF, in combination with the size of the cross-validation sample, guarantee the occurrence of near-boundary realizations. We expect that the beta-based Leja collocation will fail to produce accurate approximations at the tails of the distribution, and thus result in inferior approximation accuracies.

The results regarding the interpolation accuracy are presented in Figure 3.9. We note that the CC results are identical to the ones in Figure 3.6, since the nodes remain unaffected from the input distributions. As expected, the approximation accuracy of the beta-based Leja nodes is significantly inferior to the CC or the uniform-Leja one, the latter shown in Figure 3.6. As already mentioned, we attribute this result to the difficulty of the interpolation rule in approximating the tail of the distribution. This approximation problem must be taken into account, when surrogate-based sampling approaches are considered. For example, even if the input distributions are non-uniform, it may be computationally more efficient to construct approximations based on uniform densities and then sample them with realizations drawn from the true input PDF.

The results regarding the estimations of the first two statistical moments are presented in Figure 3.10. In this case, the Leja-based estimations are significantly superior to the CC-based ones. While both rules reach machine accuracy, the convergence of Algorithm 3.1 is many times faster, if based on the weighted Leja nodes. Therefore, while inferior in terms of approximation accuracy, weighted Leja nodes seem to be very advantageous for quadrature purposes, for the non-symmetric input PDFs considered here.

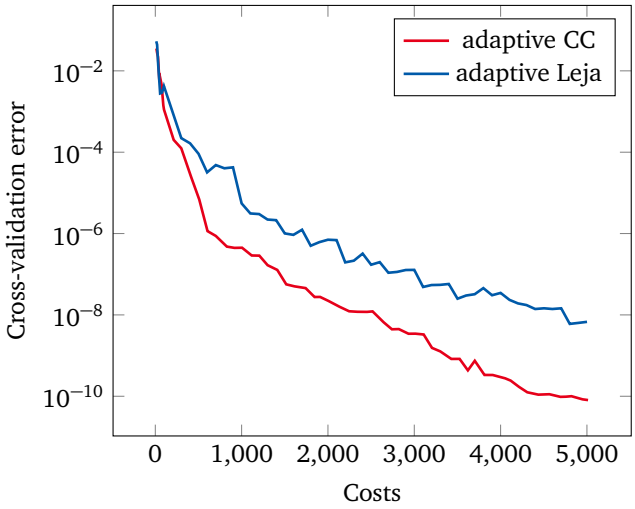

Figure 3.9: Cross-validation error versus costs for the dimension-adaptive stochastic collocation method, applied to the model waveguide problem with beta distributed random inputs.

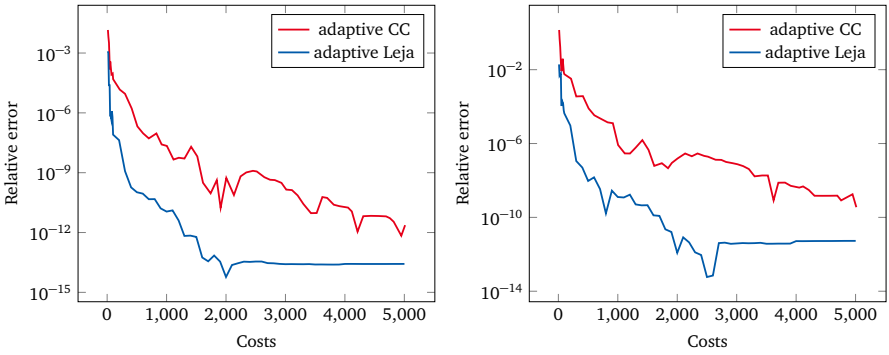

(a) Relative error versus costs for $\mathbb{E}[|S_{11}|]$.

(b) Relative error versus costs for $\mathbb{V}[|S_{11}|]$.

Figure 3.10: Moment relative errors versus costs for the dimension-adaptive stochastic collocation method, applied to the model waveguide problem with beta distributed random inputs.

3.7 Concluding Remarks

In this chapter we have introduced a dimension-adaptive stochastic collocation method and applied it for UQ purposes to a model waveguide problem. We have

verified the advantages of adaptively constructed, anisotropic sparse grids compared to isotropic Smolyak ones. For the same costs, i.e. number of model evaluations, the dimension-adaptive approach improves the approximation accuracy significantly, usually by several orders of magnitude. Equivalently, given a desired accuracy, major computational savings can be achieved using the dimension-adaptive scheme.

A central consideration of this chapter was the comparison between weighted and unweighted CC and Leja points in the context of the dimension-adaptive stochastic collocation method. We have confirmed the observations of the available literature regarding the advantages of Leja nodes for interpolation purposes, as well as their suboptimal performance for quadrature purposes, in the case of uniformly distributed inputs. The results are reversed in the case of non-uniform distributions, specifically regarding the skewed beta distributions considered here. In that case, a comparison between a weighted Leja rule and a CC rule with modified quadrature weights shows that the former is inferior in terms of approximation accuracy, however, yields significantly more accurate results in terms of quadrature.

4 Basis and Sampling-Adaptive Generalized Polynomial Chaos

In this section we present a number of algorithms for the adaptive construction of generalized polynomial chaos (gPC) approximations. We first recall the fundamentals of univariate and multivariate gPC approximations. Methods for the computation of the gPC series coefficients are presented next. Emphasis is placed on regression-based techniques, which will also be used in the adaptive algorithms. The main topic of this chapter, i.e. adaptively constructed gPC expansions based on basis- and sampling-adaptivity, follows. The derivation of statistical measures regarding the QoI, such as moments and sensitivity indices, is discussed next. Finally, the performance and accuracy of the here-presented adaptive algorithms are verified on the model waveguide problem from Section 2.1.1 and compared against isotropic gPC and an adaptive approach based on least angle regression (LAR).

4.1 Generalized Polynomial Chaos

The gPC has its roots in the work of N. Wiener on the use of Hermite polynomials for modeling Gaussian-based stochastic processes [161]. Cameron and Martin generalized the use of Hermite polynomials for arbitrary stochastic processes with a finite second moment in [32]. Hermite polynomial chaos was established in the work of Ghanem and Spanos [61] on stochastic finite elements for engineering applications. Shortly after, Xiu and Karniadakis extended polynomial chaos from Hermite polynomials only to orthogonal polynomials from the Askey scheme [4], resulting in what is now known as the generalized or Wiener-Askey polynomial chaos [164]. In the gPC context, the spectral approximation (2.24) consists of polynomials which satisfy the orthogonality property

$$\mathbb{E}\left[\Psi_i \Psi_j\right] := \int_{\Xi} \Psi_i(\mathbf{y})\Psi_j(\mathbf{y})\varrho(\mathbf{y})\mathrm{d}\mathbf{y} = \mathbb{E}\left[(\Psi_i)^2\right]\delta_{ij}, \tag{4.1}$$

as well as the properties

$$\mathbb{E}[\Psi_0] = 1, \quad \text{and} \quad \mathbb{E}[\Psi_k] = 0, k > 0. \tag{4.2}$$

We will here focus on gPC approximations based on the random discrete projection (RDP) approach [106]. In this case, we start with a set of realizations $\{\mathbf{y}_l\}_{l=1}^{L}$, typically called the experimental design. Unlike the stochastic collocation case (see Chapter 3), the realizations do not correspond to predefined interpolation nodes, but are sampled from the joint input PDF $\varrho(\mathbf{y})$ in a random or quasi-random way [25, 103, 104]. The model is evaluated on the realizations and the approximation is computed as the solution of a discrete least squares (LS) problem, as will be explained in Section 4.1.3. We consider here only unweighted discrete LS, as in [36, 40, 105, 106]. Approximations based on weighted discrete LS have been considered in [2, 41]. A survey on both approaches can be found in [42].

4.1.1 Univariate Polynomial Chaos Expansions

We first consider the case of a single RV Y with the univariate PDF $\varrho(y)$. Denoting a univariate polynomial of order $p \in \mathbb{N}_0$ with ψ_p, the orthogonality condition reads

$$\mathbb{E}\left[\psi_p \psi_q\right] := \int_{\Xi} \psi_p(y)\psi_q(y)\varrho(y)\mathrm{d}y = \mathbb{E}\left[\left(\psi_p\right)^2\right]\delta_{pq}, \tag{4.3}$$

while the properties in (4.2) are transformed into

$$\mathbb{E}[\psi_0] = 1, \quad \text{and} \quad \mathbb{E}\left[\psi_p\right] = 0, p > 0. \tag{4.4}$$

In Table 4.1, we present the orthogonal polynomials corresponding to four commonly employed continuous distributions. The Wiener-Askey scheme [164] considers similar correspondences for discrete distributions as well.

The univariate gPC expansion reads

$$g(y) \approx \widetilde{g}(y) = \sum_{p=0}^{p_{\max}} s_p \psi_p(y), \tag{4.5}$$

where $p_{\max}$ is the maximum polynomial order and $s_p \in \mathbb{R}$ are scalar coefficients. We will delay the computation of the series coefficients s_p until Section 4.1.3. We denote with $\mathbb{P}_{p_{\max}}$ the polynomial space

$$\mathbb{P}_{p_{\max}} := \operatorname{span}\left\{\psi_p : p \leq p_{\max}\right\}. \tag{4.6}$$

4.1.2 Multivariate Polynomial Chaos Expansions

Proceeding to the multivariate case, we introduce the multi-index $\mathbf{p} = (p_1, p_2, \ldots, p_N) \in \mathbb{N}_0^N$ which contains the polynomial order per parameter. The corresponding multivariate polynomial $\Psi_{\mathbf{p}}(\mathbf{y})$ is given by

$$\Psi_{\mathbf{p}} = \prod_{n=1}^{N} \psi_{p_n}(y_n). \tag{4.7}$$

Table 4.1: Correspondence between continuous distributions and orthogonal polynomials according to the Wiener-Askey scheme [164].

Distribution	Polynomials	Support
Uniform	Legendre	$[a, b]$
Normal	Hermite	$(-\infty, +\infty)$
Beta	Jacobi	$[a, b]$
Gamma	Laguerre	$[0, +\infty)$

The orthogonality condition in the multivariate case reads

$$\mathbb{E}\left[\Psi_{\mathbf{p}}\Psi_{\mathbf{q}}\right] = \int_{\Xi} \Psi_{\mathbf{p}}(\mathbf{y})\,\Psi_{\mathbf{q}}(\mathbf{y})\,\varrho(\mathbf{y})\,\mathrm{d}\mathbf{y} = \mathbb{E}\left[\Psi_{\mathbf{p}}^2\right]\delta_{\mathbf{pq}}, \tag{4.8}$$

where $\delta_{\mathbf{pq}} = \delta_{p_1q_1}\delta_{p_2q_2}\cdots\delta_{p_Nq_N}$. The properties in (4.2) are equivalently written as

$$\mathbb{E}[\Psi_{\mathbf{0}}] = 1, \quad \text{and} \quad \mathbb{E}\left[\Psi_{\mathbf{p}}\right] = 0, \mathbf{p} \neq \mathbf{0}, \tag{4.9}$$

where $\mathbf{0}$ is the zeroth multi-index, i.e. $\mathbf{0} = (0, 0, \ldots, 0)$. A simple way to construct the multivariate polynomial basis is by tensorizing the univariate bases, such that

$$\{\Psi_{\mathbf{p}}\}_{\mathbf{p}\in\Lambda^{\mathrm{TP}}_{p_{\max}}} := \{\psi_{p_1}\}_{p_1=0}^{p_{\max,1}} \otimes \{\psi_{p_2}\}_{p_2=0}^{p_{\max,2}} \otimes \cdots \otimes \{\psi_{p_N}\}_{p_N=0}^{p_{\max,N}}, \tag{4.10}$$

where $p_{\max} := \max_n p_{\max,n}$ and $\Lambda^{\mathrm{TP}}_{p_{\max}}$ is the tensor-product multi-index set defined as

$$\Lambda^{\mathrm{TP}}_{p_{\max}} := \left\{\mathbf{p} \in \mathbb{N}_0^N : \max_n p_n \leq p_{\max}\right\}. \tag{4.11}$$

Using the TP approach the cardinality of the polynomial basis increases exponentially with the number of parameters, i.e. $\#\Lambda^{\mathrm{TP}}_{p_{\max}} = p_{\max}^N$. As per its definition [162, 164], the gPC basis is a total degree (TD) polynomial basis, consisting of all polynomials defined by multi-indices in the set

$$\Lambda^{\mathrm{TD}}_{p_{\max}} := \{\mathbf{p} \in \mathbb{N}_0^N : |\mathbf{p}| := p_1 + p_2 + \cdots + p_N \leq p_{\max}\}. \tag{4.12}$$

Then, the number of basis terms is equal to $\#\Lambda^{\mathrm{TD}}_{p_{\max}} = \frac{(N+p_{\max})!}{N!p_{\max}!}$. A comparison between TP and TD bases for different numbers of parameters N is presented in Figure 4.1. Irrespective of the choice of polynomial basis, the multivariate gPC approximation reads

$$g(\mathbf{y}) \approx \widetilde{g}(\mathbf{y}) = \sum_{\mathbf{p}\in\Lambda} s_{\mathbf{p}}\Psi_{\mathbf{p}}(\mathbf{y}). \tag{4.13}$$

The corresponding multivariate polynomial space $\mathbb{P}_\Lambda$ is given by

$$\mathbb{P}_\Lambda = \mathrm{span}\left\{\Psi_{\mathbf{p}} : \mathbf{p} \in \Lambda\right\}. \tag{4.14}$$

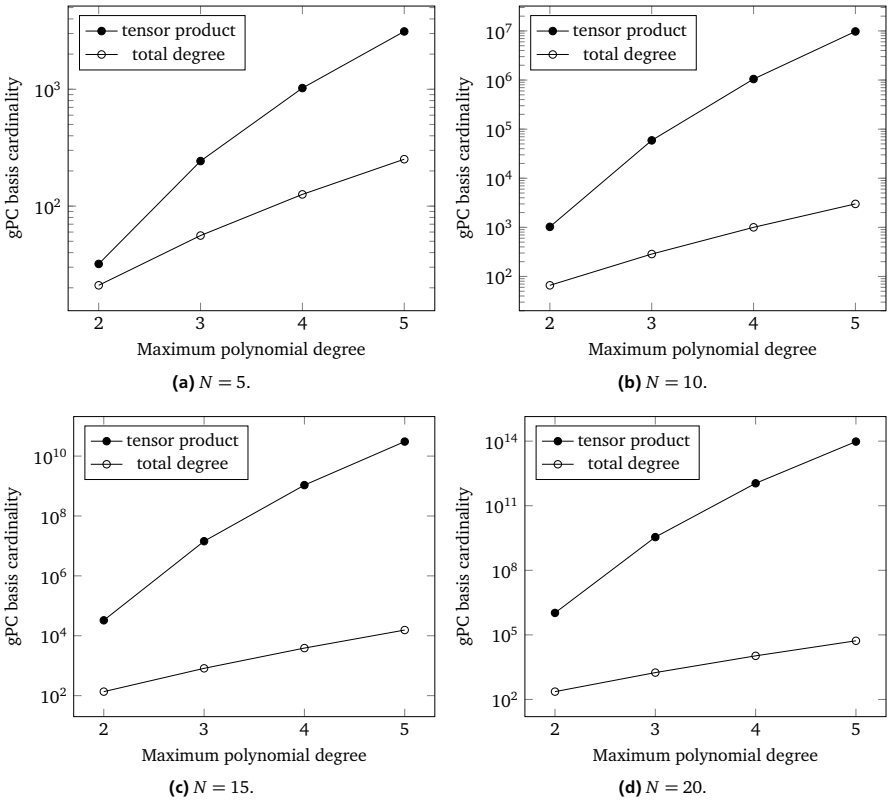

Figure 4.1: Complexity growth of tensor product and total degree gPC polynomial bases.

4.1.3 Computing the Polynomial Chaos Coefficients

The literature offers two main choices for the computation of the series coefficients $s_{\mathbf{p}}$ in the context of non-intrusive gPC approximations, namely least squares (LS) regression [22, 24], also referred to as random discrete projection (RDP) in [105, 106], and pseudo-spectral projection [49, 88]. In this work, the adaptive methods detailed in Section 4.2 are based on the LS-based computation of coefficients, therefore, this approach will be the main focus of this section. However, we also present the pseudo-spectral projection approach, for completeness. We note that the two approaches are equivalent, if infinitely many samples are con-

sidered in the regression case. We also note that we use unweighted discrete LS [36, 40, 105, 106] in this work. Regarding weighted LS approximations, the interested reader is referred to [2, 41, 42].

Pseudo-spectral projection

We first multiply (4.13) with $\Psi_{\mathbf{q}}$, where $\mathbf{q} \in \Lambda$ and Λ is either the TP or TD multi-index set. Then, we apply the expectation operator on both sides of the resulting equation. Under the orthogonality condition (4.8), the series coefficients $s_{\mathbf{p}}$ are given as

$$s_{\mathbf{p}} = \frac{\mathbb{E}\left[g\Psi_{\mathbf{p}}\right]}{\mathbb{E}\left[\Psi_{\mathbf{p}}^2\right]} = \frac{\int_{\Xi} g(\mathbf{y})\,\Psi_{\mathbf{p}}(\mathbf{y})\,\varrho(\mathbf{y})\,\mathrm{d}\mathbf{y}}{\int_{\Xi} \Psi_{\mathbf{p}}(\mathbf{y})\,\Psi_{\mathbf{p}}(\mathbf{y})\,\varrho(\mathbf{y})\,\mathrm{d}\mathbf{y}}. \tag{4.15}$$

In (4.15), the multivariate integrals of the numerator are typically computed by means of numerical integration [87], e.g. (quasi-) MC sampling or Gauss quadrature. For Wiener-Askey polynomials [164], the denominator can be determined analytically.

Least squares regression

We start by considering an experimental design, i.e. a set of realizations $\{\mathbf{y}_l\}_{l=1}^{L}$ and the corresponding set of model evaluations $\{g(\mathbf{y}_l)\}_{l=1}^{L}$. The gPC approximation is obtained as the result of the minimization problem

$$\widetilde{g} = \underset{\pi \in \mathbb{P}_{\Lambda}}{\arg\min} \sum_{l=1}^{L} \left(g(\mathbf{y}_l) - \pi(\mathbf{y}_l)\right)^2. \tag{4.16}$$

Assuming that a polynomial basis $\left\{\Psi_{\mathbf{p}}\right\}_{\mathbf{p}\in\Lambda}$ is available and has a cardinality $\#\Lambda = M$, we introduce the design matrix $\mathbf{D} \in \mathbb{R}^{L\times M}$, with elements $\mathbf{D}_{lm} = \Psi_m(\mathbf{y}_l)$, and the vector of model evaluations $\mathbf{g} = (g(\mathbf{y}_1), g(\mathbf{y}_2)\ldots, g(\mathbf{y}_L))$. Collecting the unknown gPC coefficients into a vector $\mathbf{s} = (s_1, s_2, \ldots, s_M)$, we form the minimization problem

$$\mathbf{s} = \underset{\hat{\mathbf{s}}\in\mathbb{R}^M}{\arg\min} \|\mathbf{D}\hat{\mathbf{s}} - \mathbf{g}\|_2. \tag{4.17}$$

By transforming (4.17) into the normal equation, we obtain the linear system

$$\mathbf{D}^{\top}\mathbf{D}\mathbf{s} = \mathbf{D}^{\top}\mathbf{g}, \tag{4.18}$$

where the system matrix $\mathbf{D}^{\top}\mathbf{D}$ is called the information matrix. The system is solvable with a unique solution if the information matrix is regular. This is the case

if $L \geq M$, i.e. the corresponding system of equations in (4.18) is overdetermined, and $\text{rank}(\mathbf{D}) = M$, i.e. the design matrix has full rank.

Due to conditioning issues, it is generally not recommended to directly invert the information matrix in (4.18). A Cholesky decomposition of a symmetric positive definite information matrix can be employed for the solution of (4.18), instead. Alternatively, a QR decomposition of the design matrix $\mathbf{D}$ can be employed, such that

$$\mathbf{D} = \mathbf{QR}, \quad \text{with } \mathbf{R} = \begin{pmatrix} \widetilde{\mathbf{R}} \\ \mathbf{0} \end{pmatrix} \in \mathbb{R}^{L \times M}, \tag{4.19}$$

where $\mathbf{Q} \in \mathbb{R}^{L \times L}$ is orthogonal and $\widetilde{\mathbf{R}} \in \mathbb{R}^{M \times M}$ is upper triangular. Then, (4.17) can be transformed into the equivalent problem

$$\|\mathbf{D}\hat{\mathbf{s}} - \mathbf{g}\|_2 = \|\mathbf{QR}\hat{\mathbf{s}} - \mathbf{QQ}^\top \mathbf{g}\|_2 = \|\mathbf{R}\hat{\mathbf{s}} - \mathbf{Q}^\top \mathbf{g}\|_2, \tag{4.20}$$

the solution of which is given as

$$\hat{\mathbf{s}} = \widetilde{\mathbf{R}}^{-1} \left(\widetilde{\mathbf{Q}}^\top \mathbf{g} \right), \tag{4.21}$$

where $\widetilde{\mathbf{Q}} \in \mathbb{R}^{L \times M}$ corresponds to the M first columns of $\mathbf{Q}$. The QR decomposition of the design matrix $\mathbf{D}$ is employed as solution method in this work.

In [36, 40, 105, 106], the authors present several studies regarding the stability of the (unweighted) LS problem (4.16), equivalently, the conditioning of the system in (4.18). All relevant results have also been collected in the survey paper [42]. The condition number of the information matrix $\kappa(\mathbf{D}^\top \mathbf{D})$ is used as a stability measure. In [106] it is shown that $\lim_{L \to \infty} \kappa(\mathbf{D}^\top \mathbf{D}) \to 1$, thus, the size of the experimental design is important for a well-conditioned LS problem. Assuming probability measures of the Jacobi type, e.g. uniform or Chebyshev ones, an RDP-based approximation is stable and optimally convergent under the relation $L = CM^q$, where $q \geq \ln 3 / \ln 2$ [42]. The equalities $q = \ln 3 / \ln 2$ and $q = 2$ hold for Chebyshev and uniform measures, respectively. In all other cases, it holds that $q > 2$.

Considering non-Chebyshev measures, the condition $L = CM^q$ with $q \geq 2$ has been found to be too pessimistic in practice, despite being theoretically optimal. Several numerical experiments indicate that a linear relation between the size of the experimental design and the number of basis terms is sufficient to obtain well-conditioned systems and near-optimal convergence, see e.g. [105, 106]. Therefore, a linear relation $L = CM$ is typically employed, where the constant $C > 1$ is called the oversampling coefficient. A choice of $C \in [2, 5]$ most often results in sufficiently accurate approximations, however, the extent of oversampling is generally

problem-dependent. We note that such linear relations have been theoretically proven for the stability and optimal convergence of weighted LS approximations [41, 42], however, weighted LS are not considered in this work.

In our adaptive algorithms, presented in Section 4.2, we also employ $\kappa\left(\mathbf{D}^\top\mathbf{D}\right)$ as a measure of LS stability. However, as already pointed out, we use here a QR factorization of the design matrix $\mathbf{D}$ and solve problem (4.20), in order to avoid the ill-conditioning issues related to the normal equation (4.18). Nevertheless, it is straightforward to show a direct relation between the condition numbers appearing in both cases. First, using the singular value decomposition (SVD) of the design matrix $\mathbf{D}$, it is easy to show that

$$\kappa\left(\mathbf{D}^\top\mathbf{D}\right) = \left(\kappa\left(\mathbf{D}\right)\right)^2. \tag{4.22}$$

Moreover, since $\mathbf{D} = \mathbf{QR}$ with $\mathbf{Q}$ being an orthogonal matrix, it holds that

$$\kappa(\mathbf{D}) = \kappa(\mathbf{QR}) = \kappa(\mathbf{R}). \tag{4.23}$$

We observe that the condition number of the system matrix of the normal equation (4.18) is significantly worse than the one of the design matrix and, equivalently, of the system matrix obtained with the QR decomposition.

4.2 Adaptive Polynomial Chaos Expansions

In the following, we will rely on the regression-based computation of gPC coefficients and on stability estimates in order to guide adaptivity with respect both to the construction of the gPC basis and to the expansion of the experimental design. Due to the necessary oversampling for stable LS problems, sampling adaptivity should follow basis adaptivity, therefore, we first present basis adaptivity in Section 4.2.1 and then sampling adaptivity in Section 4.2.2.

4.2.1 Basis Adaptivity

In this section, we will assume that a fixed experimental design $\{\mathbf{y}_l\}_{l=1}^L$ and the corresponding set of model evaluations $\{g\left(\mathbf{y}_l\right)\}_{l=1}^L$ are available. Accordingly, the computational costs attributed to the number of model evaluations are now fixed and equal to L. Given an oversampling coefficient C such that $L \geq CM$, where M denotes the number of approximation terms, equivalently, the size of the gPC polynomial basis, a TP or TD basis of suitable size can be constructed. However, the TP and TD bases defined by the multi-index sets $\Lambda^{\mathrm{TP}}_{p_{\max}}$ in (4.11) and $\Lambda^{\mathrm{TD}}_{p_{\max}}$ in (4.12), respectively, are isotropic, i.e. the same maximum polynomial degree is considered for every parameter. As in the dimension-adaptive stochastic collocation case, we

can exploit parameter anisotropies and construct adaptively a multi-index set Λ and the corresponding anisotropic gPC basis and approximation. Assuming that the QoI is not equally sensitive to all parameters, the anisotropic gPC expansion will be more accurate than its isotropic TP or TD counterparts for the same number of approximation terms.

Similarly to Section 3.3, we will employ a dimension-adaptive approach based on a sequence of nested, downward-closed (see Definition 2.4) multi-index sets $(\Lambda_k)_{k\geq 0}$, such that $\Lambda_0 \subset \Lambda_1 \subset \cdots \subset \Lambda_k \subset \cdots$. At every iteration, the multi-index set Λ_k is expanded by exactly one multi-index $\mathbf{p}_{k+1} \in \Lambda_k^{\text{adm}}$, such that $\Lambda_{k+1} = \Lambda_k \cup \{\mathbf{p}_{k+1}\}$. Expanding the multi-index set with multiple coefficients at each step has also been considered in [102]. The multi-index $\mathbf{p}_{k+1}$ is chosen to be the one with the maximum contribution to the already available gPC approximation among all admissible multi-indices $\mathbf{p} \in \Lambda_k^{\text{adm}}$. The contribution indicators $\eta_{\mathbf{p}}$ are now given by the absolute values of the admissible gPC coefficients, such that $\eta_{\mathbf{p}} := \left|s_{\mathbf{p}}\right|$, $\mathbf{p} \in \Lambda_k^{\text{adm}}$. The use of the gPC coefficients as contribution indicators is motivated by the fact that they are expected to decay quickly for smooth QoIs [115]. We consider two different approaches for computing the admissible gPC coefficients, detailed below.

We will refer to the first approach as the "all-in" approach, due to the fact that the gPC coefficients are computed by solving the LS problem (4.18) using a polynomial basis $\left\{\Psi_{\mathbf{p}}\right\}_{\mathbf{p}\in\Lambda_k^{\text{LS}}}$, where $\Lambda_k^{\text{LS}} = \Lambda_k \cup \Lambda_k^{\text{adm}}$, i.e. the LS problem is solved including all admissible multi-indices $\mathbf{p} \in \Lambda_k^{\text{adm}}$. Then, the multi-index $\mathbf{p}_{k+1}$ is chosen as

$$\mathbf{p}_{k+1} = \underset{\mathbf{p}\in\Lambda_k^{\text{adm}}}{\arg\max}\, \eta_{\mathbf{p}} = \underset{\mathbf{p}\in\Lambda_k^{\text{adm}}}{\arg\max} \left|s_{\mathbf{p}}\right|, \tag{4.24}$$

where $s_{\mathbf{p}}$, $\mathbf{p} \in \Lambda_k^{\text{adm}}$, are the coefficients resulting from the solution of the regression problem (4.18). This procedure is continued iteratively until a number of LS unknowns M, equivalently, a gPC basis cardinality $\#\Lambda = M$ has been reached, such that $CM \geq L$, or until the maximum or total contribution of the admissible set Λ_k^{adm} is below a specified tolerance ϵ. The adaptive basis construction with the all-in approach is shown in Algorithm 4.1. After the termination of the algorithm, the gPC approximation is constructed using the multi-index set $\Lambda = \Lambda_k \cup \Lambda_k^{\text{adm}}$ in order to include the already computed contributions of the admissible multi-indices.

The second approach will be referred to as the "one-to-one" approach, because now the LS problem (4.18) is solved once per admissible multi-index $\mathbf{p} \in \Lambda_k^{\text{adm}}$, each time using the multi-index set $\Lambda_k^{\text{LS}} = \Lambda_k \cup \{\mathbf{p}\}$, $\mathbf{p} \in \Lambda_k^{\text{adm}}$, and the corresponding polynomial basis $\left\{\Psi_{\mathbf{p}}\right\}_{\mathbf{p}\in\Lambda_k^{\text{LS}}}$. Each one of the LS solutions yields a gPC coefficient $s_{\mathbf{p}}$, $\mathbf{p} \in \Lambda_k^{\text{adm}}$. As in the all-in case, the multi-index set is expanded iteratively, such

Algorithm 4.1: Adaptive gPC polynomial basis construction with the "all-in" approach.

Data: map $g(\mathbf{y})$, initial multi-index set Λ_0, tolerance ϵ, experimental design $\{\mathbf{y}_l, g(\mathbf{y}_l)\}_{l=1}^{L}$, oversampling coefficient C

Result: multi-index set Λ, gPC polynomial basis $\left\{\Psi_{\mathbf{p}}\right\}_{\mathbf{p}\in\Lambda}$, gPC coefficients $\left\{s_{\mathbf{p}}\right\}_{\mathbf{p}\in\Lambda}$

$k = 0$
while TRUE **do**
  Compute admissible set $\Lambda_k^{\text{adm}} := \left\{\mathbf{p} \in \Lambda_k^+ : \mathbf{p} \notin \Lambda_k \quad \text{and} \quad \{\mathbf{p}\}^- \subset \Lambda_k\right\}$.
  Construct LS multi-index set $\Lambda_k^{\text{LS}} = \Lambda_k \cup \Lambda_k^{\text{adm}}$ and polynomial basis $\left\{\Psi_{\mathbf{p}}\right\}_{\mathbf{p}\in\Lambda_k^{\text{LS}}}$.
  Using $\left\{\Psi_{\mathbf{p}}\right\}_{\mathbf{p}\in\Lambda_k^{\text{LS}}}$, compute the gPC coefficients by solving $\mathbf{s} = \arg\min_{\hat{\mathbf{s}}\in\mathbb{R}^M} \|\mathbf{D}\hat{\mathbf{s}} - \mathbf{g}\|_2$.
  Compute contribution indicators $\eta_{\mathbf{p}} = \left|s_{\mathbf{p}}\right|$, $\forall \mathbf{p} \in \Lambda_k^{\text{adm}}$.
  Compute current number of LS unknowns $M_k = \#\Lambda_k^{\text{LS}} = \#\left(\Lambda_k \cup \Lambda_k^{\text{adm}}\right)$.
  Compute current total (or maximum) contribution $\epsilon_k = \sum_{\mathbf{p}\in\Lambda_k^{\text{adm}}} \eta_{\mathbf{p}}$ $\left(\text{or } \epsilon_k = \max_{\mathbf{p}\in\Lambda_k^{\text{adm}}} \eta_{\mathbf{p}}\right)$.
  if $CM_k \geq L$ OR $\epsilon_k \leq \epsilon$ **then**
    break loop
  end
  Find new multi-index $\mathbf{p}_{k+1} = \arg\max_{\mathbf{p}\in\Lambda_k^{\text{adm}}} \eta_{\mathbf{p}}$.
  Update activated set $\Lambda_{k+1} = \Lambda_k \cup \{\mathbf{p}_{k+1}\}$.
  $k = k + 1$
end
Construct the gPC polynomial basis $\left\{\Psi_{\mathbf{p}}\right\}_{\mathbf{p}\in\Lambda}$ and compute the gPC coefficients $\left\{s_{\mathbf{p}}\right\}_{\mathbf{p}\in\Lambda}$ for $\Lambda = \Lambda_k \cup \Lambda_k^{\text{adm}}$.

that $\Lambda_{k+1} = \Lambda_k \cup \{\mathbf{p}_{k+1}\}$, where $\mathbf{p}_{k+1}$ is given by (4.24), until the same termination criteria are met. The one-to-one procedure is depicted in Algorithm 4.2. Contrary to the all-in case, we cannot consider the multi-index set $\Lambda = \Lambda_k \cup \Lambda_k^{\text{adm}}$ after the termination of the algorithm, due to the fact that the size of the experimental design is not sufficient for computing the corresponding LS problem. Compared to the all-in approach, the one-to-one algorithm requires the solution of a significantly larger number of LS problems, however, of smaller size.

Algorithm 4.2: Adaptive gPC polynomial basis construction with the "one-to-one" approach.

Data: map $g(\mathbf{y})$, initial multi-index set Λ_0, tolerance ϵ, experimental design $\{\mathbf{y}_l, g(\mathbf{y}_l)\}_{l=1}^{L}$, oversampling coefficient C

Result: multi-index set Λ, gPC polynomial basis $\left\{\Psi_{\mathbf{p}}\right\}_{\mathbf{p}\in\Lambda}$, gPC coefficients $\left\{s_{\mathbf{p}}\right\}_{\mathbf{p}\in\Lambda}$

$k = 0$
while TRUE **do**
  Compute admissible set $\Lambda_k^{\text{adm}} := \left\{\mathbf{p} \in \Lambda_k^+ : \mathbf{p} \notin \Lambda_k \quad \text{and} \quad \{\mathbf{p}\}^- \subset \Lambda_k\right\}$.
  for EVERY MULTI-INDEX $\mathbf{p} \in \Lambda_k^{\text{adm}}$ **do**
    Construct LS multi-index set $\Lambda_k^{\text{LS}} = \Lambda_k \cup \{\mathbf{p}\}$ and polynomial basis $\left\{\Psi_{\mathbf{p}}\right\}_{\mathbf{p}\in\Lambda_k^{\text{LS}}}$.
    Using $\left\{\Psi_{\mathbf{p}}\right\}_{\mathbf{p}\in\Lambda_k^{\text{LS}}}$, compute the gPC coefficients by solving $\mathbf{s} = \arg\min_{\hat{\mathbf{s}}\in\mathbb{R}^M} \|\mathbf{D}\hat{\mathbf{s}} - \mathbf{g}\|_2$.
  end
  Compute contribution indicator $\eta_{\mathbf{p}} = \left|s_{\mathbf{p}}\right|$, $\forall \mathbf{p} \in \Lambda_k^{\text{adm}}$.
  Compute current number of LS unknowns $M_k = \#\Lambda_k^{\text{LS}} = \#\Lambda_k + 1$.
  Compute current total (or maximum) contribution $\epsilon_k = \sum_{\mathbf{p}\in\Lambda_k^{\text{adm}}} \eta_{\mathbf{p}}$ $\left(\text{or } \epsilon_k = \max_{\mathbf{p}\in\Lambda_k^{\text{adm}}} \eta_{\mathbf{p}}\right)$.
  if $CM_k \geq L$ OR $\epsilon_k \leq \epsilon$ **then**
    break loop
  end
  Find new multi-index $\mathbf{p}_{k+1} = \arg\max_{\mathbf{p}\in\Lambda_k^{\text{adm}}} \eta_{\mathbf{p}}$.
  Update activated set $\Lambda_{k+1} = \Lambda_k \cup \{\mathbf{p}_{k+1}\}$.
  $k = k + 1$
end
Construct the gPC polynomial basis $\left\{\Psi_{\mathbf{p}}\right\}_{\mathbf{p}\in\Lambda}$ and compute the gPC coefficients $\left\{s_{\mathbf{p}}\right\}_{\mathbf{p}\in\Lambda}$ for $\Lambda = \Lambda_k$.

4.2.2 Sampling Adaptivity

As already mentioned in Section 4.1.3, the stability of the LS problem (4.16) depends on the relation between the size of the experimental design, L, and the number of approximation terms, M, typically given as $L = CM$, where $C > 1$ is an oversampling coefficient. On the one hand, an underestimation of C will result in a

badly conditioned LS system and thus to inaccurate solutions. On the other hand, overestimating C will result in unnecessary model evaluations, equivalently, costs.

In this section, on top of the basis adaptation suggested in Section 4.2.1, we also suggest an approach for the adaptive extension of the experimental design, based on estimates regarding the stability of the LS problem. As to follow the theory presented in [36, 40, 42, 105, 106], we will use the stability estimate given by the condition number $\kappa\left(\mathbf{D}^\top\mathbf{D}\right)$. However, the equivalence to $\kappa(\mathbf{D})$ or, in the case of a QR decomposition of $\mathbf{D}$, $\kappa(\mathbf{R})$, has been shown in (4.22) and (4.23), respectively. Therefore, the same adaptive approaches can be employed using those condition numbers as well.

The sampling-adaptive procedure can be adapted with minor modifications to both the all-in and one-to-one basis adaptive methods, presented in Algorithms 4.1 and 4.2, respectively. Exemplarily, we present the all-in sampling-adaptive approach in Algorithm 4.3. In essence, we apply the basis-adaptive Algorithm 4.1 for an experimental design of fixed size, as long as the condition of the LS problem remains acceptable. For that reason, we monitor the condition number $\kappa\left(\mathbf{D}^\top\mathbf{D}\right)$. Whenever $\kappa\left(\mathbf{D}^\top\mathbf{D}\right)$ exceeds a pre-defined threshold $\kappa_{\max}$, we expand the experimental design, equivalently, we increase L, until the condition $\kappa\left(\mathbf{D}^\top\mathbf{D}\right) \leq \kappa_{\max}$ is met. The algorithm terminates if a predefined accuracy ϵ or a simulation budget B is reached.

The proposed basis and sampling adaptive approach has two main advantages. First, the condition $\kappa\left(\mathbf{D}^\top\mathbf{D}\right) \leq \kappa_{\max}$ ensures that the LS problems are stable and their solutions are accurate. Secondly, by expanding the experimental design only when required by stability constraints, we do not rely on an a priori fixed oversampling coefficient. Instead, we rely on an a priori set limit $\kappa_{\max}$, the value of which must ensure the LS stability.

The question remains as to how to choose the condition number limit $\kappa_{\max}$, i.e. which condition numbers lead to stable LS problems and thus convergent gPC approximations. In [40], the authors use the condition number limit $\kappa_{\max} = 3$, which ensures that the information matrix $\mathbf{D}^\top\mathbf{D}$ is well-conditioned with a high probability. However, similarly to the oversampling relations discussed in Section 4.1.3, the suggested condition number bound is often too pessimistic. In several numerical experiments, see e.g. [105, 106] and Section 4.4.3 in this work, it is observed that relatively large condition numbers still result in accurate gPC approximations.

In the context of this work, we have found that the constraint $\kappa\left(\mathbf{D}^\top\mathbf{D}\right) \leq 100$ is typically sufficient for converging gPC approximations (see Section 6). This can be explained due to the fact that we use the QR decomposition of the design matrix $\mathbf{D}$ to solve (4.17), in which case the condition number equivalences (4.22) and (4.23) must be taken into account. Hence, the equivalent constraint is $\kappa(\mathbf{D}) =$

$\kappa(\mathbf{R}) \leq 10$, i.e. not too large condition numbers are allowed for the actual system matrix. For the case of the simple waveguide model employed in Section 4.4, even higher values of $\kappa_{\max}$ yield accurate results. Overall, similarly to choosing a priori a suitable oversampling coefficient, we note that a good choice of $\kappa_{\max}$ is again problem-dependent, therefore, certain heuristics are employed in this work.

4.3 Post-processing the Approximation

As already mentioned in Section 2.3, once available, the gPC approximation can replace the original model in intensive sampling-based computations. However, we can also exploit the properties of the multivariate orthogonal polynomials in (4.13) to derive expected values, variances and Sobol sensitivity indices [146] by directly post-processing the gPC series terms [22, 24, 149].

We assume that the gPC approximation uses the polynomial basis $\left\{\Psi_{\mathbf{p}}\right\}_{\mathbf{p}\in\Lambda}$, the polynomials of which satisfy the properties (4.1) and (4.2), equivalently, (4.8) and (4.9). Taking advantage of those properties, the expected value $\mathbb{E}[g]$ can be estimated as

$$\begin{aligned}\mathbb{E}[g] = \int_{\Xi} g(\mathbf{y})\,\varrho(\mathbf{y})\,\mathrm{d}\mathbf{y} &\approx \int_{\Xi} \sum_{\mathbf{p}\in\Lambda} s_{\mathbf{p}}\Psi_{\mathbf{p}}(\mathbf{y})\,\varrho(\mathbf{y})\mathrm{d}\mathbf{y} \\ &= \sum_{\mathbf{p}\in\Lambda}\left(s_{\mathbf{p}}\int_{\Xi}\Psi_{\mathbf{p}}(\mathbf{y})\,\varrho(\mathbf{y})\mathrm{d}\mathbf{y}\right) \\ &= \sum_{\mathbf{p}\in\Lambda} s_{\mathbf{p}}\mathbb{E}\left[\Psi_{\mathbf{p}}\right] = s_{\mathbf{0}},\end{aligned} \tag{4.25}$$

where $\mathbf{0}$ is the zeroth multi-index, i.e. $\mathbf{0} = (0,0,\ldots,0)$. Similarly, the variance $\mathbb{V}[g]$ can be estimated as

$$\begin{aligned}\mathbb{V}[g] &= \mathbb{E}\left[(g-\mathbb{E}[g])^2\right] = \mathbb{E}\left[g^2\right] - (\mathbb{E}[g])^2 \\ &\approx \left(\int_{\Xi}\left(\sum_{\mathbf{p}\in\Lambda} s_{\mathbf{p}}\Psi_{\mathbf{p}}(\mathbf{y})\right)^2 \varrho(\mathbf{y})\mathrm{d}\mathbf{y}\right) - s_{\mathbf{0}}^2 \\ &= \left(\sum_{\mathbf{p}\in\Lambda} s_{\mathbf{p}}^2\int_{\Xi}\Psi_{\mathbf{p}}(\mathbf{y})^2\,\varrho(\mathbf{y})\mathrm{d}\mathbf{y}\right) - s_{\mathbf{0}}^2 \\ &= \left(\sum_{\mathbf{p}\in\Lambda} s_{\mathbf{p}}^2\mathbb{E}\left[\Psi_{\mathbf{p}}^2\right]\right) - s_{\mathbf{0}}^2 = \sum_{\mathbf{p}\in\Lambda\setminus\mathbf{0}} s_{\mathbf{p}}^2\mathbb{E}\left[\Psi_{\mathbf{p}}^2\right].\end{aligned} \tag{4.26}$$

Algorithm 4.3: Adaptive basis and sampling gPC approximation construction with the "all-in" approach.

Data: map $g(\mathbf{y})$, initial multi-index set Λ_0, experimental design $\{\mathbf{y}_l, g(\mathbf{y}_l)\}_{l=1}^{L}$, maximum condition number $\kappa_{\max}$, simulation budget B, tolerance ϵ

Result: multi-index set Λ, gPC polynomial basis $\left\{\Psi_{\mathbf{p}}\right\}_{\mathbf{p}\in\Lambda}$, gPC coefficients $\left\{s_{\mathbf{p}}\right\}_{\mathbf{p}\in\Lambda}$

$k = 0$
while TRUE **do**
  Compute admissible set $\Lambda_k^{\text{adm}} := \left\{\mathbf{p} \in \Lambda_k^+ : \mathbf{p} \notin \Lambda_k \quad \text{and} \quad \{\mathbf{p}\}^- \subset \Lambda_k\right\}$.
  Construct LS multi-index set $\Lambda_k^{\text{LS}} = \Lambda_k \cup \Lambda_k^{\text{adm}}$ and polynomial basis $\left\{\Psi_{\mathbf{p}}\right\}_{\mathbf{p}\in\Lambda_k^{\text{LS}}}$.
  Construct design matrix $\mathbf{D}$ and compute $\kappa\left(\mathbf{D}^\top\mathbf{D}\right)$.
  while $\kappa\left(\mathbf{D}^\top\mathbf{D}\right) > \kappa_{\max}$ **do**
    Increase experimental design size L.
    Construct design matrix $\mathbf{D}$ and compute $\kappa\left(\mathbf{D}^\top\mathbf{D}\right)$.
  end
  Using $\left\{\Psi_{\mathbf{p}}\right\}_{\mathbf{p}\in\Lambda_k^{\text{LS}}}$, compute the gPC coefficients by solving $\mathbf{s} = \arg\min_{\hat{\mathbf{s}}\in\mathbb{R}^M} \|\mathbf{D}\hat{\mathbf{s}} - \mathbf{g}\|_2$.
  Compute contribution indicators $\eta_{\mathbf{p}} = \left|s_{\mathbf{p}}\right|$, $\forall \mathbf{p} \in \Lambda_k^{\text{adm}}$.
  Compute current number of LS unknowns $M_k = \#\Lambda_k^{\text{LS}} = \#\left(\Lambda_k \cup \Lambda_k^{\text{adm}}\right)$.
  Compute current total (or maximum) contribution $\epsilon_k = \sum_{\mathbf{p}\in\Lambda_k^{\text{adm}}} \eta_{\mathbf{p}}$ $\left(\text{or } \epsilon_k = \max_{\mathbf{p}\in\Lambda_k^{\text{adm}}} \eta_{\mathbf{p}}\right)$.
  if $L \geq B$ OR $\epsilon_k \leq \epsilon$ **then**
    break loop
  end
  Find new multi-index $\mathbf{p}_{k+1} = \arg\max_{\mathbf{p}\in\Lambda_k^{\text{adm}}} \eta_{\mathbf{p}}$.
  Update activated set $\Lambda_{k+1} = \Lambda_k \cup \{\mathbf{p}_{k+1}\}$.
  $k = k + 1$
end
Construct the gPC polynomial basis $\left\{\Psi_{\mathbf{p}}\right\}_{\mathbf{p}\in\Lambda}$ and compute the gPC coefficients $\left\{s_{\mathbf{p}}\right\}_{\mathbf{p}\in\Lambda}$ for $\Lambda = \Lambda_k \cup \Lambda_k^{\text{adm}}$.

The computation of Sobol indices is based on estimations of partial variances, as presented in Section 2.2. We will focus here on the first-order (main effect) and

total-order (total effect) indices. We define the corresponding first- and total-order multi-index sets $\Lambda_n^{\mathrm{FO}} \subset \Lambda$ and $\Lambda_n^{\mathrm{TO}} \subset \Lambda$, $n = 1, 2, \ldots, N$, such that

$$\Lambda_n^{\mathrm{FO}} := \{\mathbf{p} \in \Lambda \,:\, p_n \neq 0 \text{ and } p_m = 0, m \neq n\}, \tag{4.27}$$

$$\Lambda_n^{\mathrm{TO}} := \{\mathbf{p} \in \Lambda \,:\, p_n \neq 0\}. \tag{4.28}$$

We then compute the gPC-based partial variances $\mathbb{V}_n^{\mathrm{FO}}[g]$ and $\mathbb{V}_n^{\mathrm{TO}}[g]$, such that

$$\mathbb{V}_n^{\mathrm{FO}}[g] = \sum_{\mathbf{p} \in \Lambda_n^{\mathrm{FO}}} s_{\mathbf{p}}^2 \mathbb{E}\left[\Psi_{\mathbf{p}}^2\right], \tag{4.29}$$

$$\mathbb{V}_n^{\mathrm{TO}}[g] = \sum_{\mathbf{p} \in \Lambda_n^{\mathrm{TO}}} s_{\mathbf{p}}^2 \mathbb{E}\left[\Psi_{\mathbf{p}}^2\right], \tag{4.30}$$

as also shown in [23, 149]. Then, the first- and total-order Sobol indices S_n^{FO} and S_n^{TO} defined in (2.20) and (2.22), respectively, are estimated as

$$S_n^{\mathrm{FO}} = \frac{\mathbb{V}_n^{\mathrm{FO}}[g]}{\mathbb{V}[g]}, \tag{4.31}$$

$$S_n^{\mathrm{TO}} = \frac{\mathbb{V}_n^{\mathrm{TO}}[g]}{\mathbb{V}[g]}. \tag{4.32}$$

Sensitivity indices besides first- and total-order, taking into account effects due to specific variable combinations, can be computed in a similar fashion.

4.4 Application to the Model Problem

We again use the stochastic parametric waveguide model presented in Section 2.1.1 to perform a series of numerical experiments. All waveguide parameters are assumed to follow uniform distributions $\mathcal{U}(a_n, b_n)$, where $a_n = y_n^{\mathrm{nom}} - 0.1 \cdot y_n^{\mathrm{nom}}$ and $b_n = y_n^{\mathrm{nom}} + 0.1 \cdot y_n^{\mathrm{nom}}$, where the nominal parameter values y_n^{nom} can be found in Table 2.3. However, we note that all adaptive gPC approaches presented in this chapter can be employed for any distribution in the Wiener-Askey scheme [164].

We first present the benefits of the basis-adaptive Algorithms 4.1 and 4.2 compared to total degree bases, for different choices of oversampling coefficients, in Section 4.4.1. Next, in Section 4.4.2, the two basis-adaptive schemes are compared against one another. In Section 4.4.3 we examine the impact of random and quasi-random sampling schemes upon the conditioning of the LS problems and upon the approximation accuracy. The basis and sampling adaptive Algorithm 4.3 is tested

against basis adaptivity with fixed oversampling coefficients in Section 4.4.4. In Section 4.4.5, Algorithm 4.3 is compared against an established adaptive gPC algorithm based on LAR [24, 113].

A few words regarding the implementation of the adaptive gPC methods:

- The total degree gPC approximations are computed with the OpenTURNS software [14].
- The implementations of the adaptive gPC Algorithms 4.1, 4.2, and 4.3 are part of an in-house developed Python software, partially based on OpenTURNS.
- The MATLAB toolbox UQLab [99] is used for the adaptive gPC expansions based on LAR.

Finally, we note that the use of randomly generated experimental designs results in a certain randomness regarding the gPC approximation itself, specifically regarding the values of the gPC coefficients. To address this issue, all results presented in this section refer to averaged results, obtained by 10 different gPC approximations which are constructed with 10 different random experimental designs. However, we note that the differences among the results of the 10 gPC approximations are minor, even considering the best and worst cases.

4.4.1 Adaptive versus Total Degree gPC Bases

We first compare gPC approximations based on TD and adaptively constructed bases. In the comparisons we employ three different oversampling coefficients, $C = 2, 3, 5$, each resulting in different costs, i.e. model evaluations, $L = CM$, for a gPC basis of cardinality M. The bases are compared in terms of approximation and moment estimation accuracy versus costs. Approximation accuracy is based on the cross-validation error (2.11) from a sample of size $Q_{\text{cv}} = 10^5$. Moment estimation accuracy is based on the relative error (2.13) using reference values of machine accuracy. The results are presented in Figures 4.2 and 4.3, respectively. The results shown in this section refer only to the "all-in" Algorithm 4.1. The results of the "one-to-one" Algorithm 4.2 are very similar, therefore, they have been omitted.

As can be observed from both the TD and adaptive basis results, the smallest oversampling coefficient, $C = 2$, seems to be the most beneficial one in terms of the accuracy-cost relation. Increasing the sample size, equivalently, the oversampling coefficient C, does not improve the accuracy in this test case and results only in unnecessary costs. Irrespective of the oversampling coefficient's value and concerning both approximation and moment estimation accuracy, the adaptively

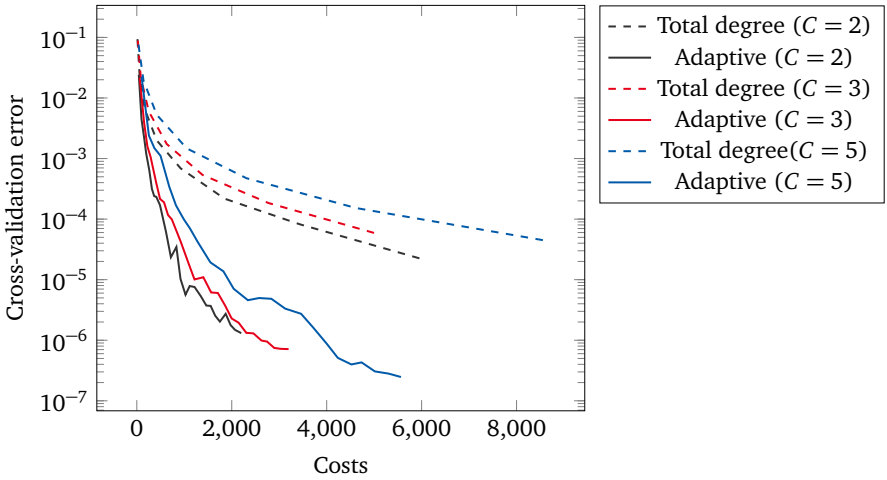

Figure 4.2: Cross-validation errors for the total degree and adaptive gPC bases and constant oversampling coefficients, applied to the model waveguide problem. The adaptive bases are constructed with the "all-in" Algorithm 4.1.

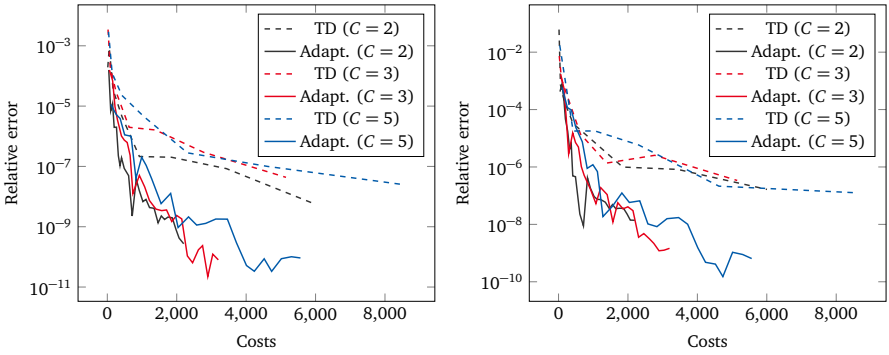

(a) Relative error versus costs for $\mathbb{E}[|S_{11}|]$.

(b) Relative error versus costs for $\mathbb{V}[|S_{11}|]$.

Figure 4.3: Moment relative errors for the total degree and adaptive gPC bases, applied to the model waveguide problem. The adaptive bases are constructed with the "all-in" Algorithm 4.1.

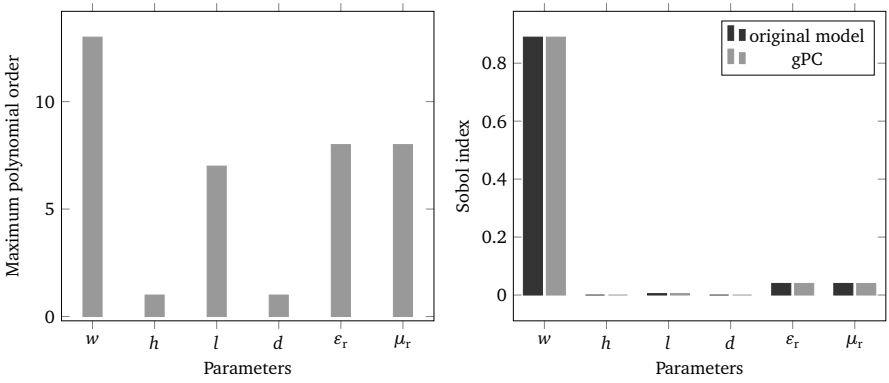

(a) Maximum polynomial order per parameter.

(b) Sobol sensitivity index per parameter.

Figure 4.4: Sensitivity analysis and adaptivity results of the adaptive gPC basis approach, applied to the model waveguide problem.

constructed bases outperform the TD ones significantly. Typically, for the same costs, error improvements of at least two orders of magnitude are obtained.

Next, we investigate the reasons behind the success of the adaptive gPC basis construction. In Figure 4.4a we plot the maximum polynomial degrees per waveguide parameter, for an adaptively constructed gPC basis with 10^3 terms. A strong anisotropy with respect to the maximum polynomial degrees can be observed. Two parameters, h and d, are completely neglected by the algorithm. This anisotropy is exactly the reason why the adaptively constructed, anisotropic gPC bases outperform the isotropic TD ones.

We also perform a gPC-based Sobol sensitivity analysis, as shown in Section 4.3, the results of which are shown in Figure 4.4b. For verification purposes, the Sobol indices are also computed using Saltelli's algorithm [135] with the original model for a sample of $Q_{\text{SA}} = 10^4$ input realizations. We only present first-order indices, since they are identical to the total-order ones. As can be observed, the gPC-based Sobol indices are identical to the sampling-based ones. We also note that the results agree with the ones given in Section 3.6.1 for the adaptive collocation method. However, in the gPC case, the Sobol indices are directly derived out of the approximation's terms and no surrogate-based sampling is necessary. As in the collocation case, with the exception of the completely insignificant parameters, the polynomial order refinement cannot be translated into sensitivity information.

4.4.2 All-In versus One-to-One Basis Adaptivity

We next compare the two proposed algorithms for the adaptive construction of the gPC basis, i.e. the "all-in" approach of Algorithm 4.1 and the "one-to-one" approach of Algorithm 4.2. We focus on approximation accuracy results alone, as the moment estimation results do not offer any further information. Using a cross-validation sample of size $Q_{\text{cv}} = 10^5$, the relation between the cross-validation error and the number of model evaluations for both algorithms is given in Figure 4.5, for oversampling coefficients $C = 2, 3, 5$. Both algorithms seem to have a similar performance, however, the one-to-one Algorithm 4.2 typically results in less model evaluations for the same accuracy.

At the same time, Algorithm 4.2 also requires solving significantly more LS problems, as already mentioned in Section 4.2.1. Exemplarily, for the case of a simulation budget of $B = 1500$ model evaluations and a fixed oversampling coefficient $C = 3$, Algorithm 4.1 solves 499 LS problems, while Algorithm 4.2 more than $3{\cdot}10^4$. Despite the fact that the LS problems solved in Algorithm 4.2 are of smaller sizes compared to the ones in Algorithm 4.2, the sheer number of LS solutions affects the overall computation time severely. Moreover, while an improved accuracy can be observed for Algorithm 4.2, the difference is nearly negligible. Therefore, we will in the following only use the "all-in" Algorithm 4.1.

4.4.3 Random versus Quasi-Random Experimental Designs

The theoretical results regarding the stability of the LS problems presented in [2, 36, 40, 41, 105, 106] refer to a design matrix $\mathbf{D} \in \mathbb{R}^{L \times M}$ constructed for a randomly generated experimental design $\{\mathbf{y}_l\}_{l=1}^{L}$. Other works, see e.g. [25, 104], suggest the use of low-discrepancy sequences instead. We consider here adaptively constructed gPC bases using Algorithm 4.1 for an oversampling coefficient $C = 2$ and experimental designs based on random and quasi-random realizations, the latter generated with Halton and Sobol sequences. We compare all approaches in terms of LS stability and approximation accuracy, the former based on the condition number $\kappa\left(\mathbf{D}^\top \mathbf{D}\right)$ and the latter on a randomly generated cross-validation sample of size $Q_{\text{cv}} = 10^5$.

The corresponding results are presented in Figure 4.6, for an increasing cardinality of the gPC basis, equivalently, for increasing costs. We observe that experimental designs based on Halton or Sobol sequences result in both more stable LS problems and improved cross-validation errors, compared to randomly generated ones. However, the available literature indicates that the advantages of low-discrepancy sequences deteriorate as the number of parameters increase [104]. In fact, quasi-random sequences are not free of the curse of dimensionality [31]. Hence, for the

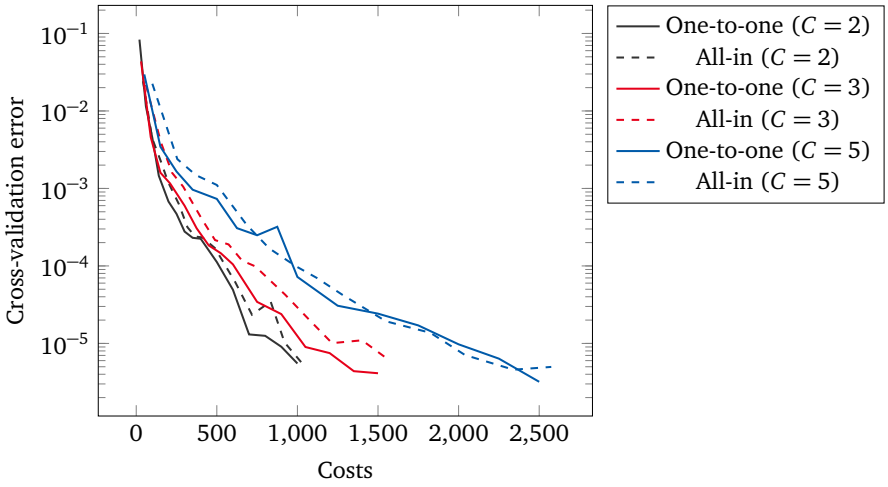

Figure 4.5: Cross-validation error versus costs for the "all-in" and "one-to-one" adaptive gPC basis approaches, applied to the model waveguide problem. The costs refer to the number of model evaluations alone.

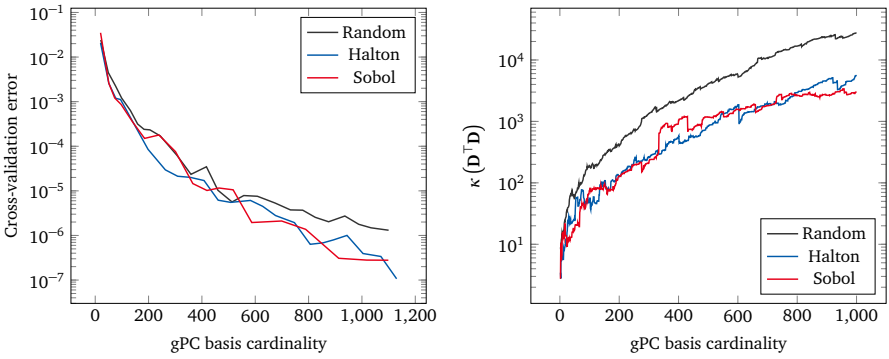

(a) Cross-validation error versus gPC basis cardinality.

(b) Condition number versus gPC basis cardinality.

Figure 4.6: Impact of different sampling schemes on gPC approximation accuracy and on LS stability.

high-dimensional models considered in Chapter 6, randomly generated experimental designs will be used.

4.4.4 Basis/Sampling Adaptivity versus Basis Adaptivity

As already pointed out in Section 4.2.2, approaches based on fixed oversampling coefficients C might result in unstable LS problems and thus inaccurate approximations or in excessive costs, for the cases of too small or too large values of C, respectively. This is also obvious in Figure 4.6b, where the condition number $\kappa\left(\mathbf{D}^\top\mathbf{D}\right)$ grows with the cardinality of the basis. While the accuracy of the gPC approximation continues to improve even for relatively large $\kappa\left(\mathbf{D}^\top\mathbf{D}\right)$ (see Figure 4.6a), stagnation or even divergence is expected after a certain point where the LS problem becomes unstable, although not observed in this particular case.

The basis and sampling-adaptive Algorithm 4.3 has been proposed as a remedy to a priori choices of C, such that the size of the experimental design is adaptively expanded according to LS stability constraints, thus avoiding both unstable LS problems and excessive model evaluations. Therefore, considering only the "all-in" strategy, we compare here Algorithm 4.3 against Algorithm 4.1. We also investigate the influence of the condition number constraint, as given by $\kappa_{\max}$ in Algorithm 4.3, on the convergence of the gPC approximation.

The results regarding the approximation accuracy for increasing costs are presented in Figure 4.7. The basis and sampling adaptive Algorithm 4.3 is employed for different condition number constraints $\kappa\left(\mathbf{D}^\top\mathbf{D}\right) \leq \kappa_{\max}$, for $\kappa_{\max} = 10^3, 10^4, 10^6$ and 10^8. We remind the reader that the LS problem (4.17) is actually solved with a QR decomposition of the design matrix $\mathbf{D}$. Following (4.22) and (4.23), the equivalent constraint is $\kappa(\mathbf{D}) = \kappa(\mathbf{R}) \leq \sqrt{\kappa_{\max}}$. The results obtained with the basis-adaptive gPC approximation with the most advantageous error-costs relation, i.e. the one for an oversampling coefficient $C = 2$, are also given in the plot, for comparison purposes.

As can be observed, a too optimistic constraint, given by $\kappa_{\max} = 10^8$, results in a poor approximation, since the LS problems are severely ill-conditioned. Lower values of $\kappa_{\max}$ result in performances comparable to the best basis-adaptive gPC approximation. Surprisingly, a good performance of Algorithm 4.3 is obtained even for $\kappa_{\max} = 10^6$. The approximation is initially relatively poor, compared to smaller values of $\kappa_{\max}$, but the performance becomes comparable for sufficiently large basis terms. However, such a result should not be generally expected. Indeed, a maximum condition number $\kappa_{\max} = 10^2$ yields the most accurate results in all high-dimensional models considered in Section 6. While no gains in terms of approximation accuracy for the same costs can be claimed, the main advantage of Algorithm 4.3 is the "on the fly" adaptation of the costs to the given LS conditioning requirements, thus avoiding a possibly poor a priori choice of the oversampling coefficient.

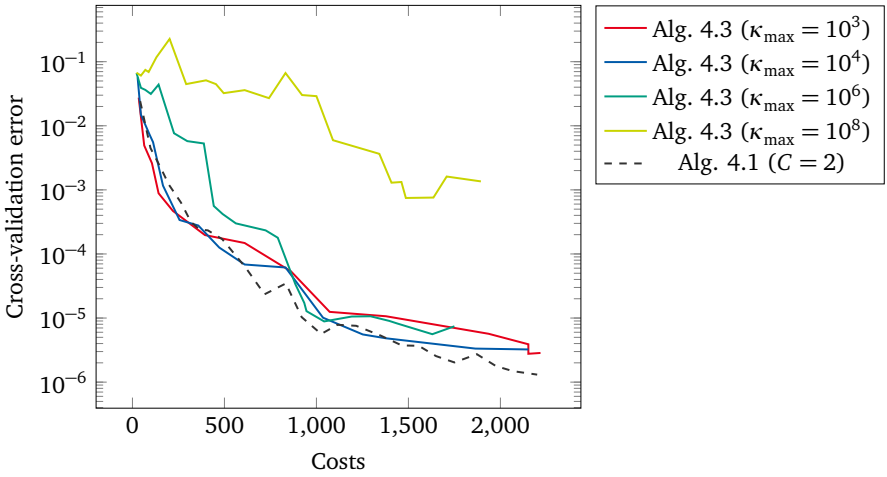

Figure 4.7: Comparison between basis/sampling-adaptive and basis-adaptive gPC approximations for the model waveguide problem. Both algorithms are based on the "all-in" basis-adaptive approach of Algorithm 4.1.

4.4.5 Basis/Sampling Adaptivity versus Least Angle Regression

A number of strategies exist in the literature regarding adaptive regression-based gPC approximations. In [102], the authors consider a fixed experimental design and propose a basis-adaptive approach based on downward-closed polynomial spaces, similar to Algorithm 4.1. A greedy approach using the orthogonal matching pursuit algorithm for selecting the gPC basis has been suggested in [78]. In [22], the authors present an iterative scheme that adds and removes basis terms according to their importance, resulting in a basis, equivalently, multi-index set, of arbitrary shape. The latter algorithm is further enriched with sequential experimental designs based on nested latin hypercube sampling (LHS), taking into account LS stability, similarly to Algorithm 4.3. The same authors have proposed a sparse, adaptive gPC method based on LAR [24], again without enforcing downward-closed multi-index sets. Finally, three algorithms for sequential sampling based on optimal weighted LS approximations are given in [2].

We focus here on fully adaptive schemes, i.e. with respect to both the gPC basis and the experimental design. Therefore, we compare our Algorithm 4.3 and the LAR-based algorithm of [24] regarding their approximation accuracy for the same costs. The cross-validation errors are computed from a set with $Q_{\text{cv}} = 10^5$ samples. The results are presented in Figure 4.8, where we can clearly see that the LAR-based

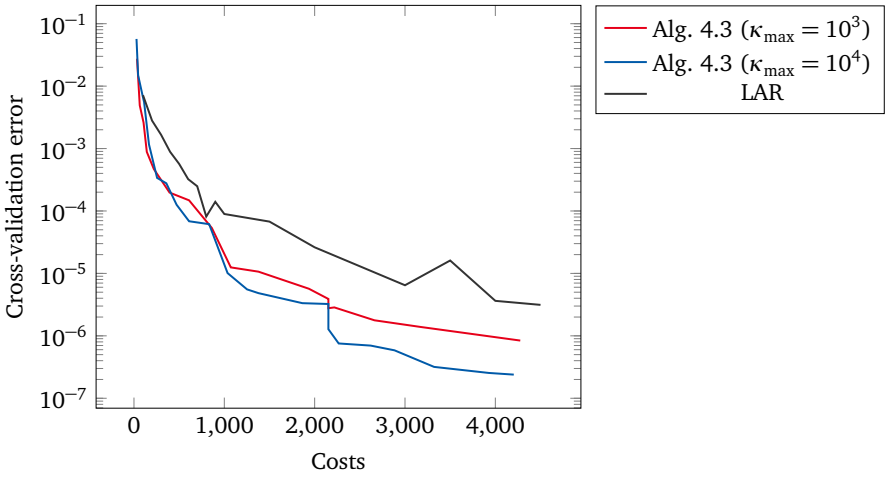

Figure 4.8: Comparison between adaptive gPC approximations based on Algorithm 4.3 and LAR.

approach is outperformed by Algorithm 4.3 for both choices of $\kappa_{\max}$. However, we note that a single numerical experiment is not sufficient for a general statement.

4.5 Concluding Remarks

In this chapter we have introduced algorithms for the adaptive expansion of gPC bases and experimental designs. Basis adaptivity is based on downward-closed multi-index sets and a dimension-adaptive approach, similar to the one presented in Chapter 3. The adaptive expansion of experimental designs takes into account requirements regarding the stability of the LS problems which must be solved to compute the gPC coefficients.

The algorithms proposed in this chapter have been compared against non-adaptive gPC approximations and against an established adaptive approach. Considering an academic waveguide toy problem, our approaches have been found to be advantageous in all cases, as better accuracies are obtained for the same costs. We note, however, that further numerical experiments must be performed for a more general statement regarding the relative performance of the compared adaptive gPC approaches.

5 Low-Rank Tensor Decompositions

In this chapter, we present a method for the computation of statistical moments based on low-rank tensor decompositions. We first present the notation and some necessary preliminaries on multi-linear algebra. Next, we present three commonly used tensor decompositions, along with their properties. We proceed by showing how tensor decompositions can be employed in the UQ context, in particular for the estimation of statistical moments. The chapter closes with the application of tensor-based UQ approaches to the model waveguide problem presented in Section 2.1.1.

5.1 Basics of Multi-Linear Algebra

In the context of the present work, a tensor $\mathcal{T}$ is an N-dimensional array, or, more formally, an element of the tensor product of N vector spaces [84], e.g. the polynomial spaces $\mathbb{P}_{\mathbf{i}} = \bigotimes_{n=1}^{N} \mathbb{P}_{i_n}$ defined in (3.7). We note that this notion of tensors is different than the one used in physics and engineering settings, e.g. the Maxwell stress tensor. The number of dimensions N is often referred to as the order of the tensor, while the dimensions are also called modes or ways. Consequently, an N-dimensional tensor is commonly encountered in the literature as a mode-N, or an N-way tensor.

In this work we will focus on real tensors, e.g. constructed by evaluating a real QoI on a tensor grid of nodes, such as the one defined in (3.4). Denoting the size of each tensor mode with J_n, $n = 1, 2, \ldots, N$, we write $\mathcal{T} \in \mathbb{R}^{J_1 \times J_2 \times \cdots \times J_N}$. A tensor element will be denoted by $\mathcal{T}_{j_1 j_2 \cdots j_N}$, for $j_n = 1, 2, \ldots, J_n$ and $n = 1, 2, \ldots, N$. Under this definition, a matrix $\mathbf{M} \in \mathbb{R}^{J_1 \times J_2}$ and a vector $\mathbf{v} \in \mathbb{R}^{J}$ are respectively equivalent to tensors of order 2 and 1. An illustration of the concept is provided in Figure 5.1.

In the following we provide some necessary definitions regarding linear and multi-linear algebraic operations. We restrict the presented definitions to the algebraic operations which are used in this work. Further definitions as well as properties of multi-linear algebra operations can be found in [69, 70, 84].

Definition 5.1 (Tensor inner product and norm)**.** *Given two tensors of equal sizes $\mathcal{S}, \mathcal{T} \in \mathbb{R}^{J_1 \times J_2 \times \cdots \times J_N}$, we define their inner product as*

$$\langle \mathcal{S}, \mathcal{T} \rangle := \sum_{j_1=1}^{J_1} \sum_{j_2=1}^{J_2} \cdots \sum_{j_N=1}^{J_N} \mathcal{S}_{j_1 j_2 \cdots j_N} \mathcal{T}_{j_1 j_2 \cdots j_N}. \tag{5.1}$$

The corresponding tensor norm, which is the analogous of the matrix Frobenius norm, is defined as

$$\|\mathcal{T}\| = \|\mathcal{T}\|_{\mathrm{F}} := \sqrt{\langle \mathcal{T}, \mathcal{T} \rangle}. \tag{5.2}$$

Definition 5.2 (Tensor unfoldings). *Let $\mathcal{T} \in \mathbb{R}^{J_1 \times J_2 \times \cdots \times J_N}$ be an N-dimensional tensor. Given an invertible bijective map*

$$\mathcal{V} : \{j_1\}_{j_1=1}^{J_1} \times \{j_2\}_{j_2=1}^{J_2} \times \cdots \times \{j_N\}_{j_N=1}^{J_N} \mapsto \{j\}_{j=1}^{J_1 J_2 \cdots J_N}, \tag{5.3}$$

we define the tensor vectorization (flattening) such that the elements of the vector $\mathbf{t} \in \mathbb{R}^{J_1 J_2 \cdots J_N}$ are given as

$$\mathbf{t}_j = \mathcal{T}_{j_1 j_2 \cdots j_N} \text{ such that } j = \mathcal{V}(j_1, j_2, \ldots, j_N). \tag{5.4}$$

Accordingly, using the invertible bijective map

$$\mathcal{M}_{(n)} : \{j_1\}_{j_1=1}^{J_1} \times \{j_2\}_{j_2=1}^{J_2} \times \cdots \times \{j_{n-1}\}_{j_{n-1}=1}^{J_{n-1}} \times \{j_{n+1}\}_{j_{n+1}=1}^{J_{n+1}} \times \cdots \times \{j_N\}_{j_N=1}^{J_N} \mapsto \{j\}_{j=1}^{J_1 J_2 \cdots J_{n-1} J_{n+1} \cdots J_N}, \tag{5.5}$$

we define the n-mode tensor matricization (unfolding) such that the elements of the matrix $\mathbf{T}_{(n)} \in \mathbb{R}^{J_n \times J_1 J_2 \cdots J_{n-1} J_{n+1} \cdots J_N}$ are given as

$$\mathbf{T}_{(n)ij} := \mathcal{T}_{j_1 j_2 \cdots j_N} \text{ such that } i = j_n, j = \mathcal{M}_{(n)}(j_1, j_2, \ldots, j_{n-1}, j_{n+1}, \ldots, j_N). \tag{5.6}$$

Given a partitioning of the dimension set $\mathcal{N} = \{1, 2, \ldots, N\}$ by two ordered sets $\mathcal{R} = \{r_1, r_2, \ldots, r_K\}$ and $\mathcal{C} = \{c_1, c_2, \ldots, c_L\}$, such that $\mathcal{N} = \mathcal{R} \cup \mathcal{C}$ and $N = K + L$, we define a general matricization, where the matrix rows are defined by the elements of the set $\mathcal{R}$ and the columns by the elements of the set $\mathcal{C}$, such that

$$\mathbf{T}_{(\mathcal{R} \times \mathcal{C})} \in \mathbb{R}^{I \times J}, \text{ where } I = \prod_{n \in \mathcal{R}} J_n, \; J = \prod_{n \in \mathcal{C}} J_n. \tag{5.7}$$

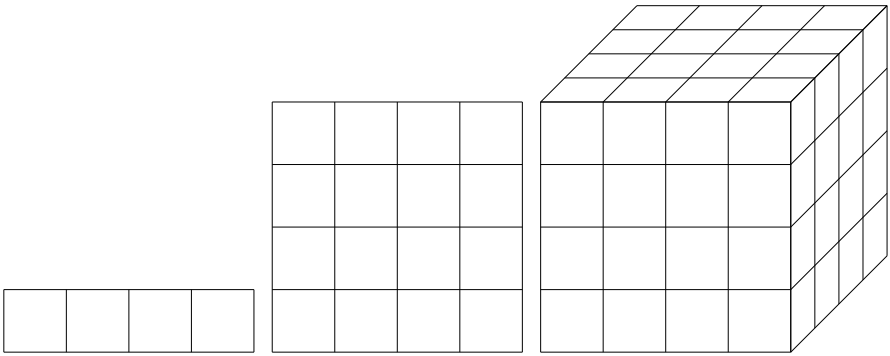

(a) Vector $\mathbf{v} \in \mathbb{R}^4$. **(b)** Matrix $\mathbf{M} \in \mathbb{R}^{4 \times 4}$. **(c)** Tensor $\mathcal{T} \in \mathbb{R}^{4 \times 4 \times 4}$.

Figure 5.1: Illustration of 1st, 2nd and 3rd-order tensors. All mode sizes are equal to 4.

Definition 5.3 (*n*-mode product)**.** *Given a tensor* $\mathcal{T} \in \mathbb{R}^{J_1 \times J_2 \times \cdots \times J_N}$ *and a matrix* $\mathbf{M} \in \mathbb{R}^{I \times J_n}$*, the n-mode (matrix) product* $\mathcal{T} \times_n \mathbf{M}$ *yields the tensor* $\mathcal{S} \in \mathbb{R}^{J_1 \times J_2 \times \cdots \times J_{n-1} \times I \times J_{n+1} \times \cdots \times J_N}$ *and is defined in elementwise notation as*

$$\mathcal{S}_{j_1 j_2 \cdots j_{n-1} i j_{n+1} \cdots j_N} = (\mathcal{T} \times_n \mathbf{M})_{j_1 j_2 \cdots j_{n-1} i j_{n+1} \cdots j_N} = \sum_{j_n=1}^{J_n} \mathcal{T}_{j_1 j_2 \cdots j_N} \mathbf{M}_{i j_n}. \tag{5.8}$$

Employing the n-mode matricization defined in (5.6)*, the n-mode product can be expressed as*

$$\mathcal{S} = \mathcal{T} \times_n \mathbf{M} \iff \mathbf{S}_{(n)} = \mathbf{M}\mathbf{T}_{(n)}. \tag{5.9}$$

Definition 5.4 (Kronecker product)**.** *Given two matrices* $\mathbf{A} \in \mathbb{R}^{I \times J}$ *and* $\mathbf{B} \in \mathbb{R}^{K \times L}$*, the Kronecker or tensor product* $\mathbf{A} \otimes \mathbf{B}$*, is a matrix* $\mathbf{C} \in \mathbb{R}^{IK \times JL}$*, such that*

$$\mathbf{C} = \mathbf{A} \otimes \mathbf{B} = \begin{bmatrix} \mathbf{A}_{11}\mathbf{B} & \mathbf{A}_{12}\mathbf{B} & \cdots & \mathbf{A}_{1J}\mathbf{B} \\ \mathbf{A}_{21}\mathbf{B} & \mathbf{A}_{22}\mathbf{B} & \cdots & \mathbf{A}_{2J}\mathbf{B} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{A}_{I1}\mathbf{B} & \mathbf{A}_{I2}\mathbf{B} & \cdots & \mathbf{A}_{IJ}\mathbf{B} \end{bmatrix}. \tag{5.10}$$

Equivalently, $\mathbf{A} \otimes \mathbf{B}$ *yields a 4th-order tensor* $\mathcal{C} \in \mathbb{R}^{I \times J \times K \times L}$*, such that* $\mathcal{C}_{ijkl} = \mathbf{A}_{ij}\mathbf{B}_{kl}$*.*

Definition 5.5 (Rank-one tensors)**.** *We call a tensor* $\mathcal{T} \in \mathbb{R}^{J_1 \times J_2 \times \cdots \times J_N}$ *elementary or rank-one if it can be written as*

$$\mathcal{T} = \mathbf{v}^{(1)} \otimes \mathbf{v}^{(2)} \otimes \cdots \otimes \mathbf{v}^{(N)}, \tag{5.11}$$

where $\mathbf{v}^{(n)} \in \mathbb{R}^{J_n}$*,* $n = 1, 2, \ldots, N$*. Equivalently to* (5.11)*, each element of a rank-one tensor* $\mathcal{T}$ *is given by*

$$\mathcal{T}_{j_1 j_2 \cdots j_N} = \mathbf{v}^{(1)}_{j_1} \mathbf{v}^{(2)}_{j_2} \cdots \mathbf{v}^{(N)}_{j_N}, \tag{5.12}$$

where $\mathbf{v}^{(n)}_{j_n}$ *is the* j_n*-th element of the vector* $\mathbf{v}^{(n)}$*.*

Definition 5.6 (Tensor rank)**.** *The rank of a tensor* $\mathcal{T} \in \mathbb{R}^{J_1 \times J_2 \times \cdots \times J_N}$ *is defined as the minimum number R of rank-one tensors that are needed for an exact representation of the tensor by means of summation, i.e.*

$$\mathcal{T} = \sum_{r=1}^{R} \mathbf{v}^{(1,r)} \otimes \mathbf{v}^{(2,r)} \otimes \cdots \otimes \mathbf{v}^{(N,r)}, \tag{5.13}$$

where $\mathbf{v}^{(n,r)} \in \mathbb{R}^{J_n}$*. We write* $\operatorname{rank}(\mathcal{T}) = R$*. For* $N = 2$*, the definition coincides with the one of the matrix rank.*

5.2 Tensor Decompositions

As can be observed, the complexity of storing all entries of tensor is $\mathcal{O}\left(J^N\right)$. The curse of dimensionality affects all tensor operations as well, such as the ones discussed above. Therefore, we will rely on tensor compression schemes, commonly called tensor decompositions, in order to avoid storing a tensor explicitly. The term "decomposition" refers to any format or scheme which allows us to express a tensor as a sequence of operations between other tensors, typically of simpler format or structure. This simpler format is most commonly based on restrictions regarding some notion of rank, as will be discussed in the following. Hence, we will refer to low-rank tensor decompositions, which can be interpreted as generalizations to low-rank matrix decompositions or factorizations.

Typically, the decomposition results in an approximation of the full tensor, $\widetilde{\mathcal{T}} \approx \mathcal{T}$, where $\widetilde{\mathcal{T}}$ is given in a format which requires less costs regarding its storage requirements and operations. In some cases, exact representations are also possible. In general, we are looking for a tensor $\widetilde{\mathcal{T}}$ which minimizes the objective function

$$f\left(\widetilde{\mathcal{T}}\right) = \|\mathcal{T} - \widetilde{\mathcal{T}}\|_{\mathrm{F}}, \tag{5.14}$$

under specific constraints regarding the format of $\widetilde{\mathcal{T}}$. The most commonly employed tensor decompositions are discussed in Sections 5.2.2, 5.2.3 and 5.2.4.

5.2.1 The Two-Dimensional Case

In the case of a matrix, equivalently, a two-dimensional tensor, $\mathbf{M} \in \mathbb{R}^{I \times J}$, a rank-$R$ decomposition reads

$$\mathbf{M} \approx \sum_{r=1}^{R} \mathbf{u}^{(r)} \otimes \mathbf{v}^{(r)}, \tag{5.15}$$

where $\mathbf{u}^{(r)} \in \mathbb{R}^I$ and $\mathbf{v}^{(r)} \in \mathbb{R}^J$. Equivalently, in elementwise format, the decomposition reads

$$\mathbf{M}_{ij} \approx \sum_{r=1}^{R} \mathbf{u}_i^{(r)} \mathbf{v}_j^{(r)} = \widetilde{\mathbf{M}}_{ij} = \sum_{r=1}^{R} \mathbf{U}_{ir} \mathbf{V}_{rj}, \tag{5.16}$$

where $\mathbf{U} \in \mathbb{R}^{I \times R}$ and $\mathbf{V} \in \mathbb{R}^{R \times J}$, with columns $\mathbf{u}^{(r)}$ and rows $\mathbf{v}^{(r)}$, respectively. In the case where $R = \operatorname{rank}(\mathbf{M})$, an exact representation in the format of (5.15) can be found.

Considering the minimization problem

$$\min_{\widetilde{\mathbf{M}}} \|\mathbf{M}-\widetilde{\mathbf{M}}\|_{\mathrm{F}} \text{ subject to } \operatorname{rank}\left(\widetilde{\mathbf{M}}\right)=R, \tag{5.17}$$

it has been shown in [51, 142] that the best rank-R matrix approximation in the Frobenius norm is obtained by a truncated singular value decomposition (SVD). However, computing the SVD of a matrix results in a high asymptotic complexity of $\mathcal{O}\left(IJ^2\right)$, assuming that $I<J$. Moreover, the full matrix must be available before the SVD is computed.

While non-optimal, a computationally more efficient alternative to the SVD is the skeleton or cross approximation [15, 26, 65, 156]. We denote the i-th row and j-th column of a matrix $\mathbf{M}\in\mathbb{R}^{I\times J}$ with $\mathbf{M}_{i:}\in\mathbb{R}^{J}$ and $\mathbf{M}_{:j}\in\mathbb{R}^{I}$, respectively. We introduce the row and column index sets $\mathscr{R}=\{i_r\}_{r=1}^{R}$ and $\mathscr{C}=\{j_r\}_{r=1}^{R}$ with cardinality $R\leq\min(I,J)$ and define the submatrices $\mathbf{M}_{\mathscr{R}:}\in\mathbb{R}^{R\times J}$, $\mathbf{M}_{:\mathscr{C}}\in\mathbb{R}^{I\times R}$, respectively holding the corresponding R rows and columns of matrix $\mathbf{M}$, as well as the intersection submatrix $\mathbf{M}_{\mathscr{R}\mathscr{C}}\in\mathbb{R}^{R\times R}$. Then, the skeleton or cross approximation of $\mathbf{M}$ is given by

$$\mathbf{M}\approx\widetilde{\mathbf{M}}=\mathbf{M}_{:\mathscr{C}}\left(\mathbf{M}_{\mathscr{R}\mathscr{C}}\right)^{-1}\mathbf{M}_{\mathscr{R}:}. \tag{5.18}$$

Computing a matrix decomposition with the cross approximation method results in an $\mathcal{O}\left((I+J)R-R^2\right)$ complexity, which is significantly lower than the one of the SVD. The main idea for computing the cross approximation can be summarized in the following procedure:

1. Start with a matrix $\mathbf{M}$, to be approximated.
2. Find a "pivot" index tuple (i^*,j^*).
3. Subtract a rank-1 cross approximation from the original matrix, such that

$$\mathbf{M}_{ij}\longleftarrow\mathbf{M}_{ij}-\frac{\mathbf{M}_{i^*j}\mathbf{M}_{ij^*}}{\mathbf{M}_{i^*j^*}} \tag{5.19}$$

4. Repeat steps (2) and (3) until the norm of $\mathbf{M}$ is below a given limit or until a pre-defined approximation rank is reached.

Different strategies for finding suitable pivot index tuples can be found in [17, 64]. The theoretically optimal index sets $\mathscr{R}$ and $\mathscr{C}$ are the ones which maximize the determinant of the intersection matrix $\mathbf{M}_{\mathscr{R}\mathscr{C}}$ over all $R\times R$ submatrices of $\mathbf{M}$ [64, 65]. However, the computation of this so-called maximum-volume submatrix is known to be an NP-hard problem. In practice, the $(I+J)R-R^2$ entries

of matrix $\mathbf{M}$ are selected in a greedy way, e.g. by using quasi maximum-volume [64, 156] or adaptive cross approximation [17, 122] algorithms. In the special case of function-generated matrices, such that $\mathbf{M}_{ij} = g(y_{1,i}, y_{2,j})$, a black-box approximation in the format of (5.18) can be computed, such that the full matrix is never explicitly formed. Those black-box cross approximations constitute the basis of the higher-order cross approximations which are employed in the context of tensor decompositions, as will be discussed next.

5.2.2 Canonical Polyadic Decomposition

The elementwise equivalent of (5.13) is given by

$$\mathcal{T}_{j_1 j_2 \cdots j_N} = \sum_{r=1}^{R} \mathbf{v}_{j_1}^{(1,r)} \mathbf{v}_{j_2}^{(2,r)} \cdots \mathbf{v}_{j_N}^{(N,r)} = \sum_{r=1}^{R} \mathbf{V}_{j_1 r}^{(1)} \mathbf{V}_{j_2 r}^{(2)} \cdots \mathbf{V}_{j_N r}^{(N)}, \tag{5.20}$$

where the vectors $\mathbf{v}^{(n,r)}$, $r = 1, 2, \ldots, R$, are the columns of the matrices $\mathbf{V}^{(n)} \in \mathbb{R}^{J_n \times R}$, called factor matrices, such that $\mathbf{V}_{j_n r_n}^{(n)} = \mathbf{v}_{j_n}^{(r,n)}$. This factorization of a tensor into a sum of elementary (rank-one) tensors constitutes the canonical polyadic (CP) decomposition, discovered independently in [33, 34, 72, 75]. In the case where the number of CP factors, R, does not coincide with the tensor rank, (5.13) and (5.20) are approximations instead of exact representations. Figure 5.2 offers an illustration for the case of a three-dimensional tensor. Normalizing the vectors $\mathbf{v}_{j_n}^{(n,r)}$ to a unit length and absorbing all normalization weights into a coefficient vector $\mathbf{c} \in \mathbb{R}^R$, the CP decomposition can be equivalently given as

$$\mathcal{T} \approx \sum_{r=1}^{R} c_r \mathbf{v}^{(1,r)} \otimes \mathbf{v}^{(2,r)} \otimes \cdots \otimes \mathbf{v}^{(N,r)} \tag{5.21}$$

$$\mathcal{T}_{j_1 j_2 \cdots j_N} \approx \sum_{r=1}^{R} c_r \mathbf{v}_{j_1}^{(1,r)} \mathbf{v}_{j_2}^{(2,r)} \cdots \mathbf{v}_{j_N}^{(N,r)} = \sum_{r=1}^{R} c_r \mathbf{V}_{j_1 r}^{(1)} \mathbf{V}_{j_2 r}^{(2)} \cdots \mathbf{V}_{j_N r}^{(N)}. \tag{5.22}$$

The CP format yields a storage complexity of $\mathcal{O}(NJR)$, which is linear with respect to the number of dimensions and, thus, free of the curse of dimensionality. A further attractive feature of the CP decomposition is its uniqueness under mild assumptions [84], which can be advantageous in practical applications. The format has been studied extensively and various algorithms have been proposed for the computation of a fixed rank-R CP decomposition, equivalently, for the solution of the minimization problem (5.14), where $\widetilde{\mathcal{T}}$ is given as in (5.21), see e.g. [21, 20]. Most algorithms rely on the alternating least squares (ALS) method [33, 72, 84], which remains the "workhorse" CP algorithm. Greedy approaches based on successive rank-1 approximations, often called generalized spectral decompositions

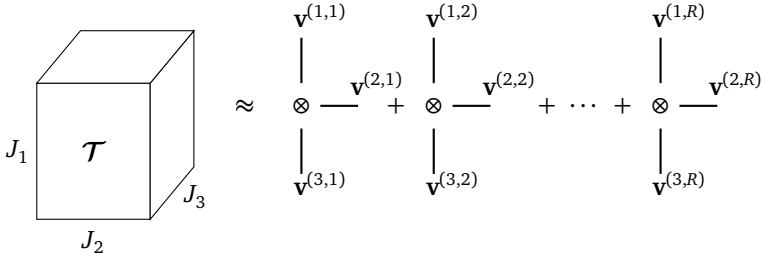

Figure 5.2: Canonical polyadic decomposition of a three-dimensional tensor.

[118], have been found to work well in the UQ context. Black-box approximations of function-generated tensors in the CP format have been investigated in [52], essentially generalizing the adaptive cross approximation idea [17, 122] beyond two dimensions.

However, the CP format faces some significant drawbacks. First, computing the rank of a tensor is an NP-hard problem [73]. More importantly, computing a fixed rank-R CP approximation generally results in an ill-posed problem [48]. It is even possible that the best rank-R CP decomposition does not exist. Therefore, the format suffers from a general lack of robustness. Accordingly, the corresponding algorithms for the numerical computation of a CP approximation for a fixed rank suffer from the same lack of robustness and may fail even in cases where good approximations are known to exist.

5.2.3 Tucker Decomposition

The Tucker format [155] decomposes a tensor $\mathcal{T} \in \mathbb{R}^{J_1 \times J_2 \times \cdots \times J_N}$ into a core tensor $\mathcal{G} \in \mathbb{R}^{R_1 \times R_2 \times \cdots \times R_N}$, where $R_n \leq J_n, \forall n = 1, 2, \ldots, N$, and N factor matrices $\mathbf{V}^{(n)} \in \mathbb{R}^{J_n \times R_n}$. The full tensor can then be approximated as

$$\mathcal{T} \approx \mathcal{G} \times_1 \mathbf{V}^{(1)} \times_2 \mathbf{V}^{(2)} \cdots \times_N \mathbf{V}^{(N)}, \tag{5.23}$$

or, in elementwise format

$$\mathcal{T}_{j_1 j_2 \cdots j_N} \approx \sum_{r_1=1}^{R_1} \sum_{r_2=1}^{R_2} \cdots \sum_{r_N=1}^{R_N} \mathcal{G}_{r_1 r_2 \cdots r_N} \mathbf{V}^{(1)}_{j_1 r_1} \mathbf{V}^{(2)}_{j_2 r_2} \cdots \mathbf{V}^{(N)}_{j_N r_N}. \tag{5.24}$$

The Tucker format can be interpreted as a multidimensional generalization of a principal component analysis [84]. An illustration for the case of a three-dimensional tensor is presented in Figure 5.3.

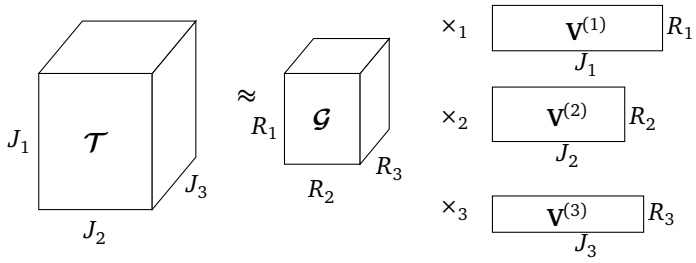

Figure 5.3: Tucker decomposition of a three-dimensional tensor.

The decomposition (5.23) is exact if $R_n = \operatorname{rank}\left(\mathbf{T}_{(n)}\right), \forall n = 1,2,\ldots,N$. In this case, the vector $\mathbf{R} = (R_1, R_2, \ldots, R_N)$ is commonly called the multi-linear or Tucker rank [46, 84]. Tucker approximations are obtained for $R_n < \operatorname{rank}\left(\mathbf{T}_{(n)}\right)$, e.g. via a truncated SVD or cross approximation of the n-mode matricization $\mathbf{T}_{(n)}$. The existence of an optimal rank-$\mathbf{R}$ Tucker approximation is guaranteed [47]. Therefore, contrary to the CP format, the Tucker format is robust. Several methods and algorithms exist for the computation of quasi-optimal rank-$\mathbf{R}$ Tucker approximations, such as the higher-order SVD [46] or the higher-order orthogonal iteration [47]. Black-box cross approximations of function-generated tensors in the Tucker format have been suggested and studied in [16, 122].

However, while robust and able to offer major computational savings if $R_n \ll J_n, \forall n = 1,2,\ldots,N$, the Tucker format still suffers from the curse of dimensionality. The number of dimensions of the core tensor $\mathcal{G}$ is equal to the number of dimensions of the original tensor $\mathcal{T}$ and the storage complexity of the Tucker format scales with $\mathcal{O}\left(NJR + R^N\right)$. Hence, for large N the format is intractable. Accordingly, the curse of dimensionality affects all algorithms which compute tensor approximations in the Tucker format.

5.2.4 Tensor-Train Decomposition

The tensor-train (TT) format [123, 125], also known as the matrix product states (MPS) format [128, 144], offers a compromise between the CP and Tucker formats, being both robust and free of the curse of dimensionality. In elementwise format, the tensor-train (TT) decomposition reads

$$\mathcal{T}_{j_1 j_2 \cdots j_N} \approx \mathbf{G}^{(1,j_1)} \mathbf{G}^{(2,j_2)} \cdots \mathbf{G}^{(N,j_N)}, \tag{5.25}$$

where $\mathbf{G}^{(n,j_n)} \in \mathbb{R}^{R_{n-1} \times R_n}$. Since the left hand side of (5.25) is a scalar value, it must hold that $R_0 = R_N = 1$. Using (5.25), every entry of the tensor is computed

by a series of matrix products, thus justifying the MPS term. An equivalent format to (5.25) which employs three-dimensional tensors $\mathcal{G}^{(n)} \in \mathbb{R}^{R_{n-1} \times J_n \times R_n}$, called TT cores, is given by

$$\mathcal{T}_{j_1 j_2 \cdots j_N} \approx \sum_{r_1=1}^{R_1} \sum_{r_2=1}^{R_2} \cdots \sum_{r_N=1}^{R_N} \prod_{n=1}^{N} \mathcal{G}^{(n)}_{r_{n-1} j_n r_n}. \tag{5.26}$$

It can easily be observed from (5.25) and (5.26) that the TT decomposition has a storage complexity $\mathcal{O}\left(NJR^2\right)$, i.e. linear with respect to the number of dimensions.

The vector $\mathbf{R} = (R_0, R_1, \ldots, R_N)$ is called the TT rank. However, contrary to the Tucker rank discussed in Section 5.2.3, the TT rank is computed by successive hierarchical SVDs, as shown in [124]. Due to this SVD hierarchy, the TT format is closely connected to the more general hierarchical Tucker format [66, 71]. A fixed rank-$\mathbf{R}$ TT decomposition can always be obtained with hierarchical truncated SVDs, hence the format is robust [125]. For computational efficiency reasons, the truncated SVDs should be avoided and replaced by more efficient matrix factorization schemes, e.g QR decompositions [125] or cross-approximations [124].

In the latter case, a black-box cross approximation of function-generated tensors in the TT format is possible. A first algorithm for that purpose has been introduced in [124]. An algorithm based on the density matrix renormalization group (DMRG) method [144, 160] is used in [137], while an algorithm based on alternating minimal energy methods is proposed in [50]. Greedy TT cross approximation algorithms are presented in [137]. All algorithms scale linearly with respect to the number of dimensions, but polynomially with respect to the TT ranks, e.g. $\mathcal{O}\left(R^2\right)$ and $\mathcal{O}\left(R^4\right)$ scalings are shown in [124] and in [50], respectively, although the hidden constants may differ significantly. Therefore, small TT ranks are crucial for low costs during a TT cross approximation.

5.3 Uncertainty Quantification with Tensor Decompositions

We consider the case of a multidimensional, tensor product (TP) quadrature scheme, similar to (3.32), used for the estimation of a statistical measure $\mathbb{E}[\phi(g)]$. For simplicity, we consider in the following the expected value of the QoI, $\mathbb{E}[g]$, however, the exact same formulation can be used for quadrature-based estimations of other statistical measures as well.

We collect the J_n univariate nodes and weights employed for each parameter in vectors $\mathbf{y}^{(n)} = \left(y_1^{(n)}, y_2^{(n)}, \ldots, y_{J_n}^{(n)}\right)$ and $\mathbf{w}^{(n)} = \left(w_1^{(n)}, w_2^{(n)}, \ldots, w_{J_n}^{(n)}\right)$, respectively. Then, the TP quadrature scheme reads

$$\mathbb{E}[g] \approx \sum_{j_1=1}^{J_1} \sum_{j_2=1}^{J_2} \cdots \sum_{j_N=1}^{J_N} g\left(\mathbf{y}_{j_1}^{(1)}, \mathbf{y}_{j_2}^{(2)}, \ldots, \mathbf{y}_{j_N}^{(N)}\right) \prod_{n=1}^{N} \mathbf{w}_{j_n}^{(n)}. \tag{5.27}$$

Next, we define the tensors $\mathcal{A} \in \mathbb{R}^{J_1 \times J_2 \times \cdots \times J_N}$ and $\mathcal{B} \in \mathbb{R}^{J_1 \times J_2 \times \cdots \times J_N}$, the elements of which are given by

$$\mathcal{A}_{j_1 j_2 \cdots j_N} = g\left(y_{j_1}^{(1)}, y_{j_2}^{(2)}, \ldots, y_{j_N}^{(N)}\right), \tag{5.28}$$

$$\mathcal{B}_{j_1 j_2 \cdots j_N} = w_{j_1}^{(1)} w_{j_2}^{(2)} \cdots w_{j_N}^{(N)}, \tag{5.29}$$

i.e. each entry of $\mathcal{A}$ corresponds to the evaluation of the QoI for a multidimensional quadrature node and each entry of $\mathcal{B}$ corresponds to a multivariate weight. Then, the quadrature formula (5.27) can be equivalently written in the form of an inner product between tensors $\mathcal{A}$ and $\mathcal{B}$, as defined in (5.1), such that

$$\mathbb{E}[g] \approx \langle \mathcal{A}, \mathcal{B} \rangle . \tag{5.30}$$

Obviously, the storage requirements and computational work of operations involving those tensors scales exponentially with the number of parameters, N. However, the curse of dimensionality can be overcome by employing the tensor decomposition methods discussed in Section 5.2.

First of all, due to its elementwise definition in (5.29), we notice that the weight tensor $\mathcal{B}$ is already given in the CP format with rank equal to 1. Moreover, the CP decomposition is exact, i.e.

$$\mathcal{B} = \mathbf{w}^{(1)} \otimes \mathbf{w}^{(2)} \otimes \cdots \otimes \mathbf{w}^{(N)}. \tag{5.31}$$

Therefore, the weight tensor's storage complexity reduces from $\mathcal{O}\left(J^N\right)$ to $\mathcal{O}(JN)$ and the curse of dimensionality is lifted. We note that (5.29) corresponds also to an exact TT decomposition with all TT ranks equal to 1.

Such a structure is not readily available in the case of the evaluations tensor $\mathcal{A}$. However, as can be seen in (5.28), $\mathcal{A}$ is a function-generated tensor, therefore, an ideal candidate for black-box tensor cross approximation methods. Assuming that a sufficiently accurate approximation can be obtained for a small CP rank R or for small multi-linear or TT ranks $\mathbf{R} = (R_1, R_2, \ldots, R_N)$, significant computational savings are expected, since the QoI will be evaluated only for the necessary tensor entries identified by the cross approximation algorithms.

5.4 Application to the Model Problem

As in Sections 3.6 and 4.4, we employ the model problem presented in Section 2.1.1 in order to verify the use of tensor decompositions for UQ purposes, as suggested above. In the following numerical investigations, all waveguide parameters are assumed to follow uniform distributions, with variations ±10% around

their nominal values, given in Table 2.3. We note that any input distribution may be used, as we are not constrained in the choice of quadrature nodes used for the tensor-product quadrature schemes. We employ Gauss-Legendre quadrature rules, which are known to be the optimal choice for the integrals considered here [153].

We first compute the expected value and the variance of the waveguide's reflection coefficient $|S_{11}|$ up to machine accuracy, using tensor-product quadrature. Machine accuracy is reached for 11 Gauss-Legendre nodes per parameter, thus, the full tensors containing the QoI evaluations on the quadrature nodes and the weights, i.e. $\mathcal{A}, \mathcal{B} \in \mathbb{R}^{J_1 \times J_2 \times \cdots \times J_6}$, $J_n = 11$, $\forall n = 1, 2, \ldots, 6$, have $11^6 = 1771561$ entries, each. The obtained moments values are used as reference values in the following.

Next, we compute tensor decompositions of the tensors corresponding to the weights and the QoI evaluations on the quadrature nodes. As already pointed out in Section 5.3 and shown in (5.31), the weights tensor $\mathcal{B}$ can be exactly represented in a rank-1 CP format, or in a rank-$(1, 1, \ldots, 1)$ TT format. A low-rank tensor approximation must be computed for the evaluations tensor $\mathcal{A}$. The approximation is also connected to the main costs of the suggested approach, since each tensor entry corresponds to a single model evaluation. Due to the format's robustness and linear complexity with respect to the number of parameters, we will rely on TT cross approximation algorithms [50, 124, 137, 138] to compute a TT approximation $\mathcal{A}_{\mathrm{TT}} \approx \mathcal{A}$ in a black-box way, i.e. without explicitly forming the full tensor. The implementations of all algorithms are available in the TT-Toolbox MATLAB software [121].

Finally, we use the TT representation of tensor $\mathcal{B}$ and the TT approximation $\mathcal{A}_{\mathrm{TT}}$ to estimate the first two moments of $|S_{11}|$, and then compute relative errors with respect to the reference values. In Figure 5.4 we present the moment estimation results obtained by applying the TT cross approximation algorithm from [50] to get $\mathcal{A}_{\mathrm{TT}}$. In the same plot, we show the TT approximation accuracy in the Frobenius norm, compared to the full tensor, i.e. the relative error $\|\mathcal{A} - \mathcal{A}_{\mathrm{TT}}\|_{\mathrm{F}} / \|\mathcal{A}\|_{\mathrm{F}}$. We note that similar results have been obtained with the algorithm from [138]. The greedy cross approximation algorithm from [137], despite being theoretically the least computationally expensive option, does not perform well for the waveguide model, stagnating at very low relative errors.

As can be observed from the results, the tensor decomposition-based moment estimations yield accurate moment estimates while using only a fraction of the full tensors. First, the weights tensor can be stored with only $NJ = 66$ entries, resulting in storage requirements less than 0.004% compared to the full tensor. Secondly, the black-box TT cross approximation of tensor $\mathcal{A}$ allows us to compute only a small number of tensor entries, compared to the full tensor. Exemplarily, the most accurate results presented in Figure 5.4 are obtained by computing approximately

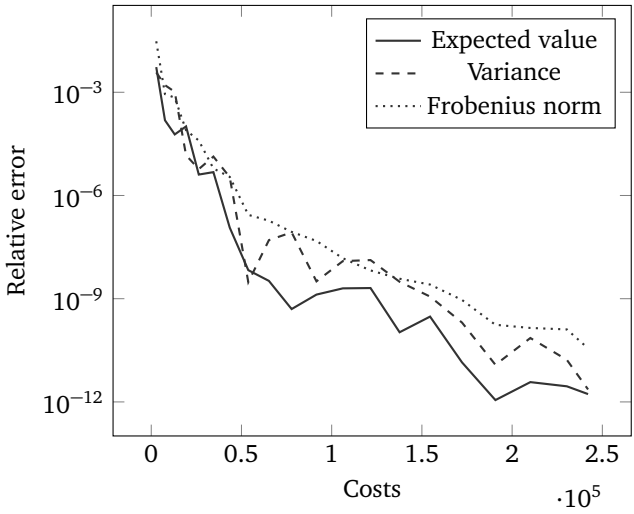

Figure 5.4: Relative errors in moment estimations and tensor approximation for the tensor decomposition-based UQ approach, applied to the model waveguide problem. The decompositions are based on the TT format. The cross approximation algorithm from [50] is employed for approximating the model evaluations tensor $\mathcal{A}$.

13% of the entries of the full tensor. It is obvious from Figure 5.4 that the accuracy of the moment estimations follows the accuracy of the TT approximation, i.e. more accurate tensor approximations yield more accurate moment estimations.

However, while indeed the considered tensor decompositions can be employed for UQ purposes and yield significant savings compared to the full tensor-product approach, the accuracy-cost ratio is significantly worse than the ones observed in Sections 3.6 and 4.4, for the dimension-adaptive stochastic collocation and the adaptive gPC methods, respectively. We note that the employed stochastic parametric waveguide model is low-dimensional, while the advantages of low-rank tensor decompositions are expected to manifest in higher-dimensional settings [157]. However, a further comparison for a model with 14 RVs, given in Section 6.1, yields similar results.

5.5 Concluding Remarks

In this chapter we presented and tested an approach based on tensor decompositions for the computation of statistical moments. The decomposition of the tensor

corresponding to the quadrature weights admits an exact rank-1 representation. A black-box cross approximation is employed for the function-generated tensor which contains the QoI evaluations on the quadrature nodes.

Using reference moment values of machine accuracy computed with a full tensor product Gauss quadrature, the decomposition-based approach is able to offer high accuracies by using only a small fraction of the full tensors. The moment estimation accuracy increases with more accurate tensor approximations of the evaluation tensor. However, compared to the corresponding results of Sections 3 and 4, the costs of the tensor-based estimations are significantly higher. We note that higher-dimensional settings might be needed to reveal the advantages of tensor decompositions, or establish a break-even against competing methods.

6 High-Dimensional Numerical Experiments

In this chapter we present three UQ studies where the respective models suffer from moderately high-dimensional input uncertainty. First, we consider an analytical permittivity model with 14 uncertain parameters. Next, we examine a Stern-Gerlach magnet model with an uncertain geometry, modeled with 14 RVs. Finally, we consider a waveguide filter, featuring 11 geometrical uncertainties. The last two studies employ relatively expensive simulation models, which, in combination with the large number of parameters, increase significantly the computational cost of the respective UQ studies. For all three cases, we present the corresponding stochastic parametric mathematical and simulation model, as well as the results of the respective UQ study.

6.1 Cole-Cole Permittivity

The electrical properties of dispersive materials are often modeled by relaxation models, which are able to capture their dispersive behavior over a wide frequency range. For example, in the case of biological tissues, the permittivity is typically modeled by either the Debye [98, 109], or the Cole-Cole [56, 134] relaxation models. The parameters of those permittivity models are usually computed by fitting them to the available experimental data [55]. Tolerances in the measurement process result in uncertainties in the experimental data, with the corresponding variations being relatively large, e.g. variations of up to 40% around the parameters' nominal values are considered in [141]. Consequently, the propagation of the uncertainty from the model parameters to the computed permittivity values must be quantified.

In this section, we will focus on a relaxation model which is commonly used in the literature to model the complex permittivity of the 17 dispersive biological tissues which form the human head, namely the fourth-order Cole-Cole model (see, e.g., [56]). Considering one RV per model parameter, the corresponding stochastic parametric model features a total of 14 RVs. Thus, although the model is an analytical one, standard approaches, such as isotropic sparse grids or total degree gPC bases, result in very high computational times.

Therefore, we employ the adaptive methods presented in Sections 3, 4 and 5 in order to compute statistical moments of the real and complex permittivity. The methods are compared against one another in terms of estimation accuracy and computational effort. Moreover, the surrogate models based on the dimension-adaptive stochastic collocation method and on the adaptive gPC approach are compared in terms of approximation accuracy. Finally, surrogate-based sensitivity analyses are presented.

6.1.1 Model

We consider the case of a biological tissue, where the dielectric spectrum is most commonly characterized by three main relaxation regions, corresponding to low, medium and high frequencies, respectively [56]. Typically, minor dispersions which do not fall into those three main regions must also be taken into account. For each relaxation region and for a given angular frequency ω, the complex relative permittivity can be computed by evaluating the Cole-Cole model

$$\varepsilon_r = \varepsilon_\infty + \frac{\varepsilon_s - \varepsilon_\infty}{1 + (\jmath\omega\tau)^{1-\alpha}} = \varepsilon_\infty + \frac{\Delta\varepsilon}{1 + (\jmath\omega\tau)^{1-\alpha}}, \tag{6.1}$$

where τ is a relaxation time constant, ε_∞ is the high-frequency relative permittivity value, i.e. for $\omega\tau \gg 1$, ε_s is the static relative permittivity value, i.e. for $\omega\tau \ll 1$, α is a measure of dispersion broadening, and $\jmath$ is the imaginary unit. The dispersion magnitude is given by the difference $\Delta\varepsilon := \varepsilon_s - \varepsilon_\infty$. We follow [56] and consider the fourth-order Cole-Cole permittivity model, which takes into account four relaxation regions and a conductivity term, such that

$$\varepsilon_r = \frac{\sigma_i}{\jmath\omega\varepsilon_0} + \varepsilon_\infty + \sum_{p=1}^{4} \frac{\Delta\varepsilon_p}{1 + \left(\jmath\omega\tau_p\right)^{1-\alpha_p}}, \tag{6.2}$$

where ε_0 is the permittivity in vacuum and σ_i denotes the static ionic conductivity. Thus, the model employs a total of 14 parameters, i.e. σ_i, ε_∞, $\Delta\varepsilon_p$, τ_p α_p, $p = 1, \ldots, 4$.

Regarding the dispersive material, we consider the case of white matter brain tissue. For this tissue, the nominal parameter values of the fourth-order Cole-Cole model have been identified in [56] and are here presented in Table 6.1. We allow a variation of ±20% around the parameters' nominal values and model them as uniformly distributed RVs, such that $Y_n \sim \mathcal{U}(a_n, b_n)$, $n = 1, 2, \ldots, N$, where $a_n = y_n^{\text{nom}} - 0.2 \cdot y_n^{\text{nom}}$ and $b_n = y_n^{\text{nom}} + 0.2 \cdot y_n^{\text{nom}}$. The ±20% variation is chosen in order to take into account relatively large uncertainties. The uniform distribution is chosen due to the lack of any further information regarding the behavior of the random parameters. For a random outcome $\theta \in \Theta$ and a fixed frequency ω, the stochastic parametric Cole-Cole model reads

$$\varepsilon_r(\theta) = \frac{\sigma_i(\theta)}{\jmath\omega\varepsilon_0} + \varepsilon_\infty(\theta) + \sum_{p=1}^{4} \frac{\Delta\varepsilon_p(\theta)}{1 + \left(\jmath\omega\tau_p(\theta)\right)^{1-\alpha_p(\theta)}}. \tag{6.3}$$

On a standard machine, a single simulation with the analytical fourth-order Cole-Cole model requires 2.5 μs, on average.

Table 6.1: Nominal parameter values for the 4th-order Cole-Cole model, corresponding to the white matter brain tissue.

Parameter	Nominal Value	Units
ε_∞	4.0	–
$\Delta\varepsilon_1$	32.0	–
$\Delta\varepsilon_2$	100.0	–
$\Delta\varepsilon_3$	$4 \cdot 10^4$	–
$\Delta\varepsilon_4$	$3.5 \cdot 10^7$	–
τ_1	7.96	ps
τ_2	7.96	ns
τ_3	53.05	μs
τ_4	7.958	ms
σ_i	0.02	S/m
α_1	0.1	–
α_2	0.1	–
α_3	0.3	–
α_4	0.02	–

6.1.2 Numerical Results

In the following, the frequency is fixed at $f = 2$ GHz. For the given frequency, we want to quantify the influence of the input RVs on the complex permittivity, therefore we monitor two QoIs, namely $g_1 = \mathrm{Re}\{\varepsilon_\mathrm{r}\}$ and $g_2 = \mathrm{Im}\{\varepsilon_\mathrm{r}\}$. We note here that all results refer to the chosen fixed frequency and changes are expected to occur when other frequency areas are considered. Ideally, a UQ study considering broadband frequencies would be desirable. However, the main goal of this study is to investigate whether the adaptive methods can be successfully applied in a moderately high-dimensional setting, as well as to offer comparisons among them in terms of accuracy and costs. Therefore, more extensive studies over a frequency range are not considered here.

We begin with estimations for the expected value and the variance of the QoIs with three different approaches. First, the estimates are obtained by post-processing the terms of a stochastic collocation approximation computed with the dimension-adaptive Algorithm 3.1. The algorithm is based on CC rules, however, the results using Leja nodes are almost identical. Secondly, the estimates are obtained by post-processing the terms of a gPC expansion computed with the basis and sampling-adaptive scheme described in Algorithm 4.3, with $\kappa_{\max} = \kappa_{\max}\left(\mathbf{D}^\top\mathbf{D}\right) = 10^2$. Thirdly, the estimates are computed using low-rank

TT approximations of the tensors corresponding to a tensor-product quadrature scheme with 6 Gauss-Legendre nodes and weights per parameter. As already shown in Section 5.3, the weights tensor admits an exact rank-$(1, 1, \ldots, 1)$ TT representation. The greedy TT cross algorithm from [137] is employed to approximate the QoI evaluations tensor. Similar results are obtained with the cross approximation algorithms from [50, 138], albeit at elevated costs for the same accuracies. We note that more quadrature nodes per parameter, equivalently, larger tensor mode sizes, do not yield any improvement in the accuracy-cost ratio.

For all approaches, we compute relative errors based on reference values obtained from a MC sampling with 10^9 random samples. The reference expected values, $\mathbb{E}[\mathrm{Re}\{\varepsilon_\mathrm{r}\}]$ and $\mathbb{E}[\mathrm{Im}\{\varepsilon_\mathrm{r}\}]$ and variances, $\mathbb{V}[\mathrm{Re}\{\varepsilon_\mathrm{r}\}]$ and $\mathbb{V}[\mathrm{Im}\{\varepsilon_\mathrm{r}\}]$, as well as the corresponding NRMSDs, are

$$\begin{aligned}
&\mathbb{E}[\mathrm{Re}\{\varepsilon_\mathrm{r}\}] = 36.993482, && \mathbb{E}[\mathrm{Im}\{\varepsilon_\mathrm{r}\}] = -9.333579,\\
&\mathbb{V}[\mathrm{Re}\{\varepsilon_\mathrm{r}\}] = 17.212063, && \mathbb{V}[\mathrm{Im}\{\varepsilon_\mathrm{r}\}] = 2.700782,\\
&\mathrm{NRMSD}[\mathrm{Re}\{\varepsilon_\mathrm{r}\}] \approx 3.6 \cdot 10^{-6}, && \mathrm{NRMSD}[\mathrm{Im}\{\varepsilon_\mathrm{r}\}] \approx 5.6 \cdot 10^{-6}.
\end{aligned}$$

For both QoIs, the relation between the relative error and the number of model evaluations, i.e. the costs, are presented in Figures 6.1, 6.2, 6.3, for the collocation, gPC and TT approaches, respectively.

All three approaches eventually stagnate at a value close to the corresponding NRMSD for both QoIs. The adaptive gPC approach stagnates after approximately

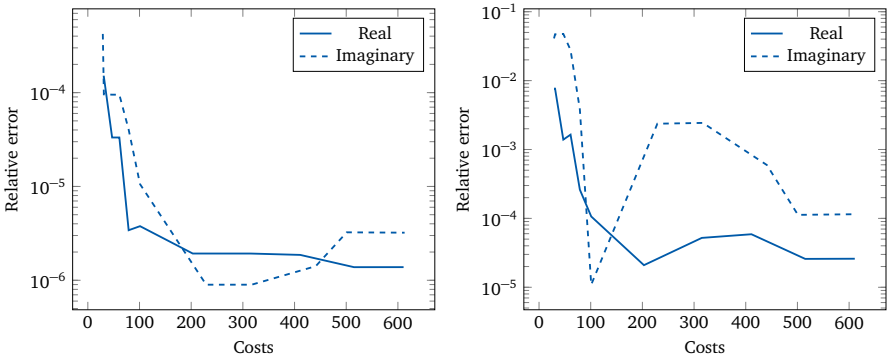

(a) Relative errors for $\mathbb{E}[\varepsilon_\mathrm{r}]$.

(b) Relative errors for $\mathbb{V}[\varepsilon_\mathrm{r}]$.

Figure 6.1: Moment relative errors versus costs for the fourth order Cole-Cole permittivity model using the dimension-adaptive stochastic collocation.

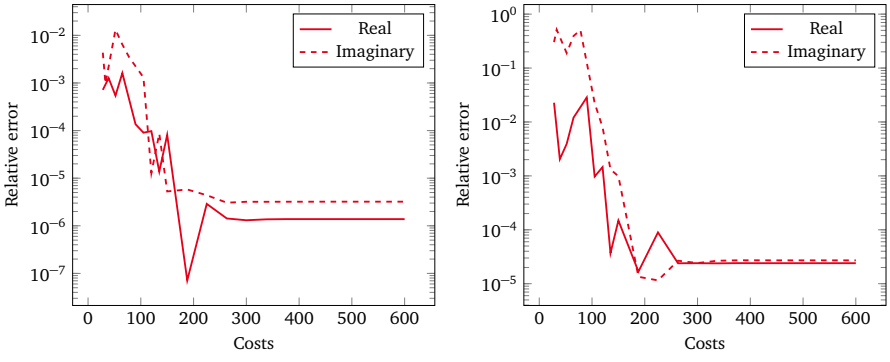

(a) Relative errors for $\mathbb{E}[\varepsilon_r]$. **(b)** Relative errors for $\mathbb{V}[\varepsilon_r]$.

Figure 6.2: Moment relative errors versus costs for the fourth order Cole-Cole permittivity model using the basis/sampling-adaptive gPC.

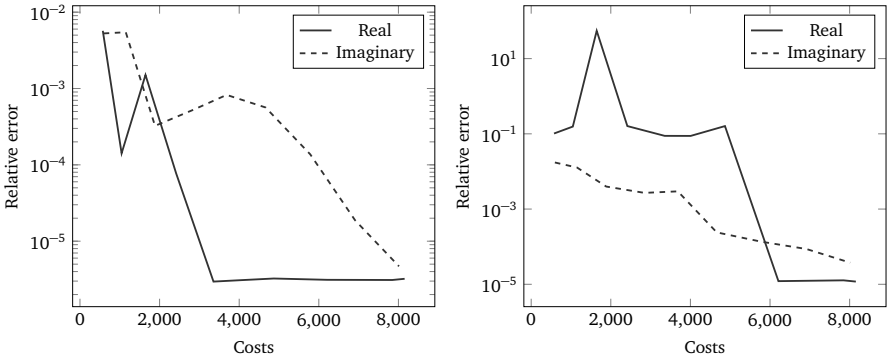

(a) Relative errors for $\mathbb{E}[\varepsilon_r]$. **(b)** Relative errors for $\mathbb{V}[\varepsilon_r]$.

Figure 6.3: Moment relative errors versus costs for the fourth order Cole-Cole permittivity model using the greedy TT cross approximation.

250 model evaluations. The collocation method needs approximately double that number of model evaluations until stagnation. For those two approaches, the same number of model evaluations leads to an error stagnation regarding the variance, as well. The tensor approximation-based approach performs noticeably worse. In the case of the real part of the permittivity, stagnation is observed after approximately 3500 model evaluations with respect to the expected value and after 6000 model

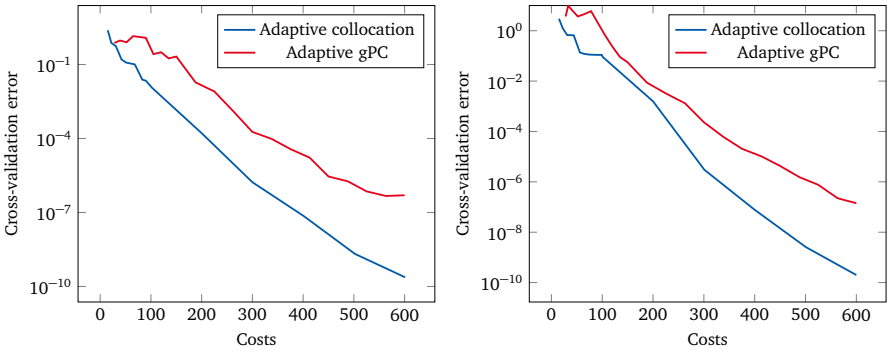

(a) Cross-validation errors versus costs for Re $\{\varepsilon_r\}$.

(b) Cross-validation errors versus costs for Im $\{\varepsilon_r\}$.

Figure 6.4: Cross-validation errors versus costs for the 4th-order Cole-Cole permittivity model. The approximation methods are the dimension-adaptive stochastic collocation and the basis/sampling-adaptive gPC.

evaluations with respect to the variance. In the case of the imaginary part, the relative errors stagnate after more than 8000 model evaluations. While the tensor compressions are massive compared to the full tensors, which would hold $6^{14} \approx 7.8 \cdot 10^{10}$ entries, the costs of the tensor-based approach exceed the costs of the other two methods by orders.

Similarly to the model waveguide problem, we observe that adaptive collocation and adaptive gPC methods are superior to tensor decompositions, at least for the here considered models and numbers of RVs. A break-even might be reached for a greater number of RVs. However, since in Sections 6.2 and 6.3 we consider comparable numbers of RVs, as well as computationally expensive numerical models instead of analytical functions, tensor decompositions will not be further employed.

We proceed with comparing the accuracy of the collocation and gPC-based surrogate models, using a cross-validation set of size $Q_{\text{cv}} = 10^5$. The results are presented in Figure 6.4. For both QoIs, the dimension-adaptive stochastic collocation has a significant advantage over the adaptive gPC. This advantage becomes more pronounced with increasing costs, at least for the number of model evaluations considered here. For the most part, the difference in the cross-validation errors of both methods for the same costs is more than two orders of magnitude.

Finally, we perform a Sobol sensitivity analysis, using the original model, the collocation-based surrogate model, and the gPC-based surrogate model. In the gPC case, the Sobol indices are directly computed by post-processing the terms

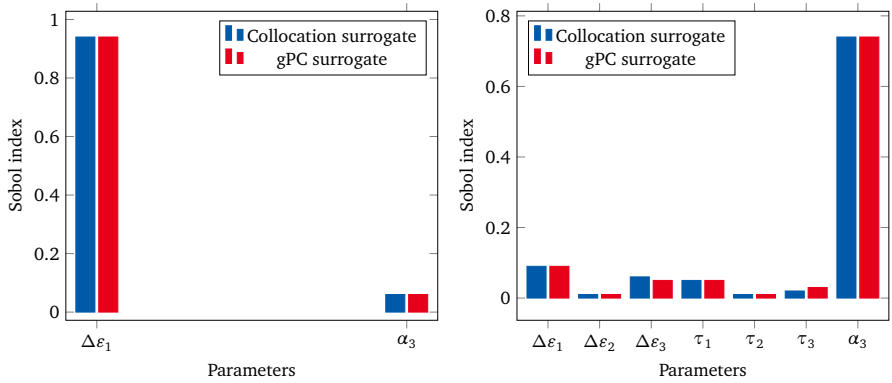

(a) Sobol sensitivity for Re $\{\varepsilon_r\}$.

(b) Sobol sensitivity for Im $\{\varepsilon_r\}$.

Figure 6.5: Sensitivity analysis results for the 4th order Cole-Cole permittivity model. The approximation methods are the dimension-adaptive stochastic collocation and the basis/sampling-adaptive gPC.

of the expansion, as shown in Section 4.3. Saltelli's algorithm [135] is used for computing the Sobol indices by sampling the original model and the collocation-based surrogate. The set of input realizations contains $Q_{\text{SA}} = 10^4$ random samples, resulting in a total of $(2N+2)Q_{\text{SA}} = 3 \cdot 10^5$ model evaluations. The results of the surrogate-based sensitivity analyses are presented in Figure 6.5. We omit the results of the original model, as they are almost identical to the surrogate-based results. Moreover, we only present first-order sensitivity indices, defined in (2.20), because their sum is almost equal to 1, hence higher order parameter interactions are negligible. We emphasize once more that the sensitivity results refer to the fixed frequency $f = 2$ GHz. Therefore, result interpretations should also take this specific setting into account.

From Figure 6.5a it can be observed that only two parameters seem to have an impact on the real part of the permittivity, namely $\Delta\varepsilon_1$ and α_3. Of those two parameters, $\Delta\varepsilon_1$ is found to be responsible for approximately 94% of the output variance, with the rest 6% being attributed to α_3. The remaining 12 parameters have negligible contributions, the sum of which does not exceed 1%. However, different sensitivities are observed in the case of the second QoI, i.e. the imaginary part of the permittivity. As shown in Figure 6.5b, parameter α_3 is now responsible for approximately 74% of the output variance, while $\Delta\varepsilon_1$ accounts for approximately 9%. Non-negligible contributions are also identified for parameters $\Delta\varepsilon_2$ (1%), $\Delta\varepsilon_3$ (6%), τ_1 (6%), τ_2 (1%), and τ_3 (3%). Consequently, regarding the complex per-

mittivity at frequency $f = 2$ GHz, at least 7 parameters could be neglected in the stochastic modeling.

6.2 Stern-Gerlach Magnet

Stern-Gerlach magnets are employed for the magnetic separation of atom beams or clusters. A key design requirement is a magnetic field with a homogeneous and strong gradient. Due to design and manufacturing imperfections, the pole region of the magnet might suffer from geometrical uncertainties, which in turn affect the field's gradient and its homogeneity. The aim of this study is to quantify the impact of geometrical uncertainties onto the average magnetic field gradient in the magnet's beam area.

We consider here the Rabi-type Stern-Gerlach magnet presented in Figure 6.6a. The magnet is similar to the one described in [100] and further studied in [45, 127, 133]. Magnets as this one are used for the deflection of nanoclusters according to their magnetic moment and are used in the setups of Stern-Gerlach experiments [58, 59]. The examined model corresponds to a real-world application, as the actual magnet is currently installed and used for various studies in KU Leuven, Belgium.

We employ the Leja-based version of Algorithm 3.1 and the basis/sampling-adaptive gPC Algorithm 4.3 to approximate the dependency of the QoI on the input parameters and derive statistical measures using the corresponding approximations. We note that using Algorithm 3.1 with CC nodes yields similar results, as shown in our work [96]. The same approaches are used to compute the expected value and the variance of the QoI. Finally, both methods are employed in a sensitivity analysis, with the aim of identifying significant and insignificant parameter contributions. The results of the sensitivity analysis show that a lower-dimensional model can be considered, such that the costs of future parameter studies, i.e. optimization, UQ, or other, may be significantly reduced. We note that the content of this section is partially based on our contribution [96, Section 4].

6.2.1 Model

All computations are performed using a linearized two-dimensional model of the magnet's cross-section, as in [127]. The only domain which is spatially resolved is the magnet's pole region, denoted with Ω_p. Domain Ω_p is decomposed into three distinct patches, namely $\Omega_\mathrm{p}^\mathrm{left}$, $\Omega_\mathrm{p}^\mathrm{air}$ and $\Omega_\mathrm{p}^\mathrm{right}$, such that $\Omega_\mathrm{p} = \Omega_\mathrm{p}^\mathrm{left} \cup \Omega_\mathrm{p}^\mathrm{air} \cup \Omega_\mathrm{p}^\mathrm{right}$. Region $\Omega_\mathrm{p}^\mathrm{air}$ refers to the air gap inside the magnet's pole region, while regions $\Omega_\mathrm{p}^\mathrm{left}$ and $\Omega_\mathrm{p}^\mathrm{right}$ to the regions on the left and right of the air gap, respectively, as shown in Figure 6.6. The magnet's beam area lies inside the air gap and is denoted with $\Omega_\mathrm{p}^\mathrm{beam}$.

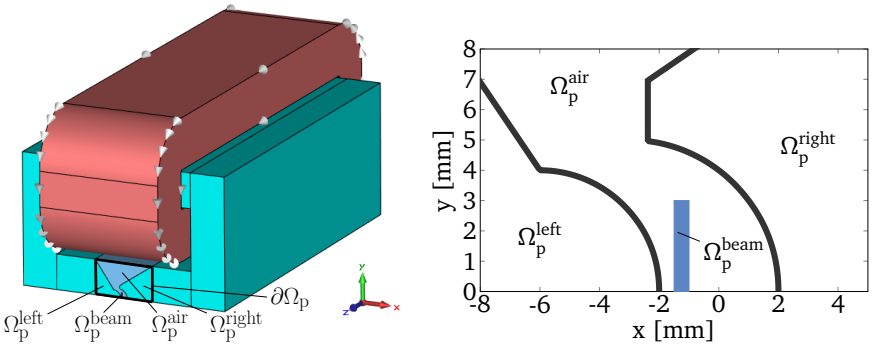

(a) 3D magnet model (half geometry).

(b) Zoom in the magnet's pole region.

Figure 6.6: The 3D model of a Rabi-type Stern-Gerlach magnet, generated with CST [44], and a zoom in the magnet's pole region. The pole region is presented for its nominal geometry. Modified pictures from [96, 127].

The remaining yoke part and the coils are modeled with a magnetic equivalent circuit and their contributions are taken into account by a field-circuit coupling approach [127]. First, the magnetic vector potential and the magnetic flux through the iron yoke are computed for the entire geometry. The values are denoted with $A_z^0(\mathbf{r})$ and Φ^0, respectively, where $\mathbf{r}$ denotes the spatial coordinate vector. The field-circuit coupling is realized by imposing

$$A_z(\mathbf{r}) = A_\Phi(\mathbf{r}) = \frac{\Phi}{\Phi^0} A_z^0(\mathbf{r}), \quad \text{on } \partial\Omega_p, \tag{6.4}$$

where Φ is recomputed for a different geometry using magnetic circuit theory. Denoting with I_c the coil current and with N_c the number of turns in the winding, the relation in (6.4) can be abstractly written as $F(A_z, \Phi) = N_c I$, where F refers to the magnetomotive force. For details, the reader is referred to [127]. In summary, the field-circuit coupled problem reads

$$\operatorname{div}(\nu \operatorname{grad} A_z) = 0, \quad \text{in } \Omega_p, \tag{6.5a}$$
$$A_z - A_\Phi = 0, \quad \text{on } \partial\Omega_p, \tag{6.5b}$$
$$F(A_z, \Phi) = N_c I_c, \tag{6.5c}$$

where $\nu := \mu^{-1}$ refers to the magnetic reluctivity. The magnetic flux density $\mathbf{B}$ is obtained as $\mathbf{B} = (\partial_y A_z, -\partial_x A_z, 0)$. Denoting with $\tau(x,y) = \frac{\partial |\mathbf{B}|}{\partial x}$ the magnetic field gradient in the x-direction, the average field gradient in the beam area is given by

$$\tau_{\text{avg}} = \frac{1}{|\Omega_{\text{p}}^{\text{beam}}|} \int_{\Omega_{\text{p}}^{\text{beam}}} \tau(x,y)\,\mathrm{d}\Omega, \tag{6.6}$$

measured in T/m.

Isogeometric analysis (IGA) [76] is employed for the spatial discretization of the pole region. In IGA, both the solution variable A_z and the geometry are described in terms of non-uniform rational B-splines (NURBS). A generic NURBS curve reads

$$\mathbf{R}(\xi) = \sum_{i=1}^{N} \mathbf{P}_i N_i^p(\xi), \quad \xi \in [0,1], \tag{6.7}$$

where $\mathbf{P}_i$ and N_i^p refer to a control point and a NURBS basis function of degree p, respectively. A NURBS basis function is defined as

$$N_i^p(\xi) = \frac{w_i B_i^p(\xi)}{\sum_{j=1}^{N} w_j B_j^p(\xi)}, \tag{6.8}$$

with weights w_i and B-spline basis functions B_i^p, respectively. The original NURBS curves defining the three patches, $\Omega_{\text{p}}^{\text{left}}$, $\Omega_{\text{p}}^{\text{air}}$, and $\Omega_{\text{p}}^{\text{right}}$, are depicted in Figure 6.6b in black.

We introduce random geometry deformations in the pole area by regarding the control points and weights of the NURBS curves as uncertain. More precisely, we introduce a total of 14 RVs, where 10 RVs correspond to the x and y coordinates of 5 control points, while 4 RVs correspond to 4 weights. The realizations of all uncertain parameters are given by $y_n = y_n^{\text{nom}} + Y_n(\theta)$, where $Y_n \sim \mathcal{U}(-1,1)$ in the case of the coordinates, and $Y_n \sim \mathcal{U}(0,1)$ for the weights. For all parameters, y_n^{nom} refers to the parameter value in the nominal pole geometry. In this way, all coordinate parameters are allowed a maximum deviation of 1mm. The random weights introduce randomness in the shape of the NURBS curves. The nominal parameter values are reported in Table 6.2.

Then, we obtain a random reluctivity as

$$\nu(\mathbf{y}) = \nu_{\text{iron}} \mathbb{1}_{\Omega_{\text{p}}^{\text{left}}(\mathbf{y})} + \nu_{\text{air}} \mathbb{1}_{\Omega_{\text{p}}^{\text{air}}(\mathbf{y})} + \nu_{\text{iron}} \mathbb{1}_{\Omega_{\text{p}}^{\text{right}}(\mathbf{y})}, \tag{6.9}$$

with $\mathbb{1}_{\Omega_p^{(\cdot)}}$ denoting the characteristic function of the corresponding patch, and ν_{iron} and ν_{air} denoting the reluctivity of iron and air, respectively. Accordingly, the parametric field-circuit coupled problem reads

$$\operatorname{div}(\nu(\mathbf{y})\operatorname{grad} A_z) = 0, \quad \text{in } \Omega_p, \tag{6.10a}$$

$$A_z - A_\Phi = 0, \quad \text{on } \partial\Omega_p, \tag{6.10b}$$

$$F(A_z, \Phi) = N_c I_c. \tag{6.10c}$$

Assuming well-posedness of the problem, the solution variable A_z is now a RV. The QoI is chosen to be the average field gradient τ_{avg}, defined in (6.6) and post-processed after A_z, such that $g = \tau_{\text{avg}}$. On a standard desktop machine, a single evaluation of the Stern-Gerlach magnet model requires approximately 60 s.

6.2.2 Numerical Results

We employ the dimension-adaptive Leja collocation and the adaptive gPC for two different maximum condition numbers, i.e. $\kappa_{\max} = 10^2$ and $\kappa_{\max} = 10^3$. We first compute relative errors with respect to the expected value and the variance of the average magnetic field gradient τ_{avg}, presented in the Figures 6.7a and 6.7b. The reference moment values are computed with a MC sampling method, using a

Table 6.2: Nominal parameter values for the Stern-Gerlach magnet.

Parameter	Nominal Value	Units
x_1	−2.38	mm
y_1	6.96	mm
x_2	−2.38	mm
y_2	4.96	mm
x_3	17.0	mm
y_3	20.0	mm
x_4	−17.0	mm
y_4	20.0	mm
x_5	−6.0	mm
y_5	4.0	mm
w_1	0.85	–
w_2	0.85	–
w_3	0.87	–
w_4	0.87	–

sample of size $Q_{\mathrm{MC}} = 10^4$. The MC-based expected value and variance, along with the corresponding RMSD and NRMSD, are

$$\begin{aligned}
\mathbb{E}\left[\tau_{\mathrm{avg}}\right] &= -238.178 \text{ T/m}, \\
\mathbb{V}\left[\tau_{\mathrm{avg}}\right] &= 21.712 \text{ (T/m)}^2, \\
\mathrm{RMSD}\left[\tau_{\mathrm{avg}}\right] &\approx 0.05 \text{ T/m}, \\
\mathrm{NRMSD}\left[\tau_{\mathrm{avg}}\right] &\approx 2 \cdot 10^{-4}.
\end{aligned}$$

Regarding the expected value, all methods converge to a relative error similar to the NRMSD given above. The adaptive Leja collocation method seems to have a slight edge over the two adaptive gPC approaches. Moreover, the adaptive gPC approach based on the smaller condition number, converges faster than the one with the larger condition number. However, the differences are rather small and all methods can be considered to perform equivalently.

In the case of the variance, again all methods converge to the same relative error. However, in this case, the adaptive gPC methods converge faster than the adaptive Leja collocation. As observed previously, a smaller condition number is advantageous for the adaptive gPC method. As in the case of the expected value, only minor differences can be observed between the gPC and the collocation-based results, hence, the overall performance of both methods in terms of moment estimation accuracy can be regarded as equivalent.

The same MC sample is used for the computation of cross-validation errors, presented in Figure 6.7c. As in the moment case, the adaptive gPC method which is based on the condition number $\kappa_{\max} = 10^2$ outperforms the one with $\kappa_{\max} = 10^3$. Both gPC approaches are initially outperformed by the adaptive Leja collocation, but are able to reach higher accuracies after more model evaluations. The difference is, however, rather small. The Leja collocation can be seen as significantly advantageous for accuracies up to 10^{-1}. For accuracies below that threshold the adaptive gPC based on $\kappa_{\max} = 10^2$ prevails slightly.

Finally, we use all surrogate models to perform a Sobol sensitivity analysis. Since the results of both gPC approaches are almost identical, we present only the ones referring to $\kappa_{\max} = 10^2$. The gPC-based Sobol indices are directly computed by post-processing the expansion terms, as shown in Section 4.3. Saltelli's algorithm [135] is used for computing the Sobol indices by sampling the collocation-based surrogate. The set of input realizations contains $Q_{\mathrm{SA}} = 10^4$ random samples, resulting in a total of $(2N+2)\,Q_{\mathrm{SA}} = 3 \cdot 10^5$ model evaluations. We remind the reader that a single solver call for the Stern-Gerlach magnet model is run in approximately 60 s, therefore, the computational costs for a sampling-based sensitivity analysis can be reduced significantly using the surrogate models.

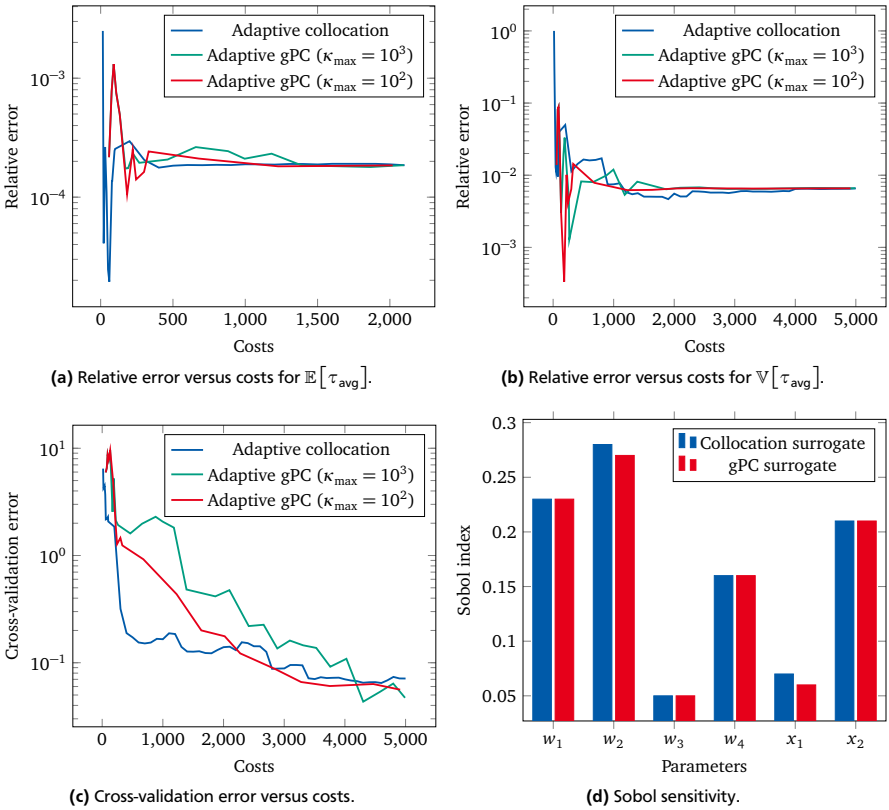

(a) Relative error versus costs for $\mathbb{E}[\tau_{\text{avg}}]$.

(b) Relative error versus costs for $\mathbb{V}[\tau_{\text{avg}}]$.

(c) Cross-validation error versus costs.

(d) Sobol sensitivity.

Figure 6.7: Uncertainty quantification and approximation results for the Stern-Gerlach magnet model. The approximation methods are the Leja-based dimension-adaptive stochastic collocation and the basis/sampling-adaptive gPC.

The results of the surrogate-based sensitivity analyses are presented in Figure 6.7d. We only present first-order sensitivity indices, defined in (2.20), because their sum is almost equal to 1, hence higher order parameter interactions are negligible. Moreover, we omit the parameters with insignificant contributions, i.e. below 1%. As can be observed, both methods recognize the same set of parameters as important ones, with only minor deviations in the Sobol index values. These minor differences can be attributed to the different sensitivity analysis methods used with each surrogate. In essence, only 6 out of the 14 parameters seem to have

an influence on the QoI, i.e. on τ_{avg}. Approximately 72% of the impact can be attributed to shape variations, given by the 4 random weights. The remaining 28% is attributed to two random coordinates, in particular the x-coordinates of points 1 and 2, with the latter having a significant contribution of 21%.

6.3 Resonant Cavity Filter

For our final numerical study, we present a case where polynomial approximations are not able to accurately reproduce the input-output dependency. We consider a real-world microwave device, in particular a Chebyshev filter constructed with waveguide technology. The basic component of the filter is a rectangular cavity, which is loaded with metallic (PEC) parts. The filter model is illustrated in Figure 6.8. Each of the seven pairs of metallic sheets is called an iris. The metallic cylinders of small radii placed in the middle of the irises' holes represent tuning screws. The metallic cylinders of larger radii placed in between irises along the z-axis are called stubs. The stubs and screws are used to tune the resonant frequencies in the cavity, e.g. by altering their position, length, or radius.

We denote the middle iris and tuning screw with "M". The filter is then separated to three segments on each side of "M" along the z-direction, each consisting of a stub, a tuning screw and an iris. Moreover, the segments are symmetric with respect to the middle "M", i.e. the geometrical characteristics of the corresponding irises, screws, and stubs are identical. The numbering of the segments, taking into account the aforementioned symmetry, is also presented in Figure 6.8. While not natural in a UQ setting, the model has originally been developed for optimization purposes, where this symmetry is necessary. Due to the symmetry, the input and output waveguide ports are interchangeable.

We consider geometrical uncertainties regarding the metallic components of the filter and try to quantify their impact on the waveguide's reflection coefficient $|S_{11}|$, which is the QoI in this study. To that end, we employ the Leja-based dimension-adaptive stochastic collocation Algorithm 3.1 and the basis/sampling-adaptive gPC Algorithm 4.3. The random variations in the geometry of the metallic components cause resonance frequency shifts, thus resulting in a non-smooth input-to-output map. As expected, the resulting surrogate models are not sufficiently accurate, as shown in Section 6.3.2. Nevertheless, despite the observed approximation problems, the moment estimation results are comparable to the ones computed with a MC sampling approach. Moreover, surprisingly, the sensitivity analysis results of both methods are found to be equivalent.

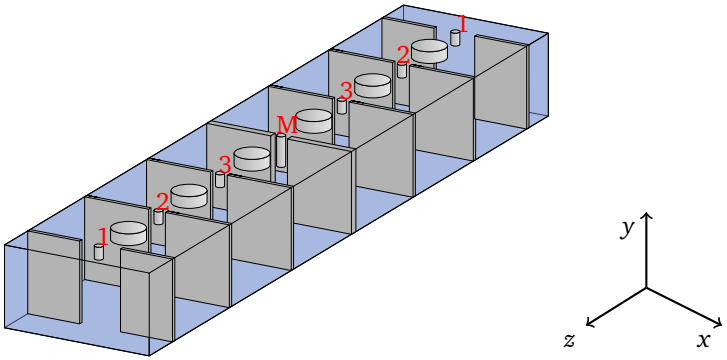

Figure 6.8: 3D model of the resonant cavity filter. The middle iris and tuning screw of the waveguide are denoted with "M". The numbers 1,2,3 correspond to the symmetric waveguide segments, each consisting of an iris, a tuning screw and a stub.

6.3.1 Model

All computations are performed using a three-dimensional FE model of the filter. The frequency-domain FE solver of the commercial CST software [44] is employed. The operation frequency is set to 670 MHz. The CST software settings allow us to consider an excitation only by the dominant mode. Under this condition, the mathematical model is given by the Maxwell's source problem (2.5), presented in Section 2.1.1.

Moving to the discrete setting, we first employ a projection on (2.5) with a test function $\mathbf{E}'$. We consider functions $\mathbf{E}, \mathbf{E}' \in \mathbf{H}_{\text{curl}}(\Omega)$, where $\mathbf{H}_{\text{curl}}(\Omega)$ denotes the space of square-integrable functions with square-integrable curl. We refer the interested reader to [107] for more information on function spaces regarding Maxwell's equations. Applying the projection and integration by parts, we get

$$\int_\Omega \mu^{-1}(\nabla \times \mathbf{E}) \cdot \left(\nabla \times \mathbf{E}'\right) \mathrm{d}\Omega - \int_\Omega \omega^2 \varepsilon \mathbf{E} \cdot \mathbf{E}' \mathrm{d}\Omega = \int_\Gamma \pi_{\mathrm{t}}\left(\mu^{-1}\nabla \times \mathbf{E}\right) \cdot \pi_{\mathrm{T}}\left(\mathbf{E}'\right) \mathrm{d}\Gamma, \tag{6.11}$$

where $\pi_{\mathrm{t}}(\cdot)$ and $\pi_{\mathrm{T}}(\cdot)$ are trace operators given by

$$\pi_{\mathrm{t}}(\mathbf{u}) = \mathbf{n} \times \mathbf{u}|_\Gamma, \tag{6.12}$$

$$\pi_{\mathrm{T}}(\mathbf{u}) = \mathbf{n} \times (\mathbf{n} \times \mathbf{u}|_\Gamma). \tag{6.13}$$

We demand the test function $\mathbf{E}'$ to fulfil the PEC boundary condition in (2.5), which results in a vanishing integral on Γ_{PEC} in the right-hand side (RHS) of (6.11). Further, we insert the port boundary conditions from (2.5) and get

$$\int_\Omega \mu^{-1} (\nabla \times \mathbf{E}) \cdot \left(\nabla \times \mathbf{E}'\right) \mathrm{d}\Omega - \int_\Omega \omega^2 \varepsilon \mathbf{E} \cdot \mathbf{E}' \mathrm{d}\Omega + \int_{\Gamma_\mathrm{P}} \jmath k_z \mu^{-1} \pi_\mathrm{T}(\mathbf{E}) \cdot \pi_\mathrm{T}\left(\mathbf{E}'\right) \mathrm{d}\Gamma_\mathrm{P} = \int_{\Gamma_\mathrm{in}} \mu^{-1} \mathbf{U}^\mathrm{inc} \cdot \pi_\mathrm{T}\left(\mathbf{E}'\right) \mathrm{d}\Gamma_\mathrm{in}, \tag{6.14}$$

where $\Gamma_\mathrm{P} = \Gamma_\mathrm{in} \cup \Gamma_\mathrm{out}$. The weak formulation then reads

$$\text{Find } \mathbf{E} \in \mathbf{V} = \left\{\mathbf{e} \in \mathbf{H}_\mathrm{curl}(\Omega) : \mathbf{n} \times \mathbf{e}|_{\Gamma_\mathrm{PEC}} = 0 \wedge \pi_\mathrm{T}(\mathbf{e}|_{\Gamma_\mathrm{P}}) \in \mathbf{L}^2(\Gamma_\mathrm{P})\right\}$$

such that

$$a(\mathbf{E}, \mathbf{E}') = \ell(\mathbf{E}'), \quad \forall \mathbf{E}' \in \mathbf{V}, \tag{6.15}$$

with

$$a(\mathbf{E}, \mathbf{E}') = \int_\Omega \mu^{-1} (\nabla \times \mathbf{E}) \cdot \left(\nabla \times \mathbf{E}'\right) \mathrm{d}\Omega - \int_\Omega \omega^2 \varepsilon \mathbf{E} \cdot \mathbf{E}' \mathrm{d}\Omega + \int_{\Gamma_\mathrm{P}} \jmath k_z \mu^{-1} \pi_\mathrm{T}(\mathbf{E}) \cdot \pi_\mathrm{T}\left(\mathbf{E}'\right) \mathrm{d}\Gamma_\mathrm{P}, \tag{6.16}$$

$$\ell(\mathbf{E}') = \int_{\Gamma_\mathrm{in}} \mu^{-1} \mathbf{U}^\mathrm{inc} \cdot \pi_\mathrm{T}\left(\mathbf{E}'\right) \mathrm{d}\Gamma_\mathrm{in}. \tag{6.17}$$

We apply the Galerkin approach, i.e. the ansatz and test functions are chosen to be the same. The electric field is approximated as

$$\mathbf{E}_h = \sum_{i=1}^{N_\mathrm{DoF}} e_i \mathbf{N}_i, \tag{6.18}$$

where e_i are the degrees of freedom (DoF), N_DoF are the number of DoF, the subscript h denotes that $\mathbf{E}$ lives in a finite function space $\mathbf{V}_h \subset \mathbf{V}$ and $\mathbf{N}_i$ denotes Nédélec basis functions of the first kind and of third order [107, 111, 166] defined on a tetrahedral mesh of the computational domain Ω. The discrete variational problem then reads

$$\text{Find } \mathbf{E}_h \in \mathbf{V}_h : \quad a(\mathbf{E}_h, \mathbf{E}_h') = \ell(\mathbf{E}_h') \quad \forall \mathbf{E}_h' \in \mathbf{V}_h. \tag{6.19}$$

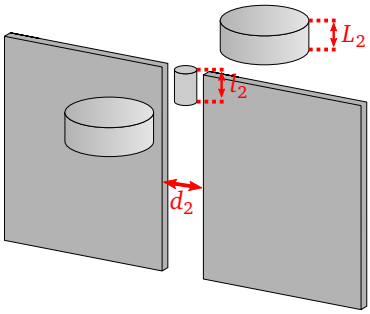

Figure 6.9: Uncertainty modeling of a single segment of the resonant cavity filter. The uncertain parameters of the (left of "M") segment 2 are shown.

Inserting (6.18) and additionally the test functions into (6.19), we obtain the linear system

$$\mathbf{A}\mathbf{e} = \mathbf{f}, \tag{6.20}$$

where $\mathbf{A} \in \mathbb{C}^{N_{\text{DoF}} \times N_{\text{DoF}}}$ is the system matrix and $\mathbf{f} \in \mathbb{C}^{N_{\text{DoF}}}$ is the discretized RHS. To obtain the reflection and/or transmission coefficients, we solve the discretized form of (2.6) with the solution of (6.20). We note that for a desired FEM accuracy of $0.5 \cdot 10^{-3}$, the FEM solver's execution time is approximately 40 minutes, using a standard desktop machine.

We consider uncertainties which do not break the segment symmetry, shown in Figure 6.8. Hence, a random outcome regarding the parameters of segment i, affects in the exact same way both symmetrical segments, thus the symmetry is retained. In particular, considering the i-th segment, where $i = 1, 2, 3$, we introduce random variations in the lengths of the stub and the screw, L_i and l_i, respectively, as well as with respect to the distance between the iris' metallic sheets, d_i. An illustration is presented in Figure 6.9. Random variations are also considered for the length of the middle screw, l_{M}, and the distance between the middle iris' sheets, d_{M}, thus resulting in a total of 11 RVs. All RVs are assumed to be uniformly distributed in the ranges given by $y_n^{\text{nom}} \pm 0.006 y_n^{\text{nom}}$, where y_n^{nom} denotes the parameter's value in the nominal geometry. The nominal parameter values are presented in Table 6.3. The maximum variation of ±0.6% with respect to the nominal values corresponds to actual manufacturing tolerances.

The introduction of the RVs results in a parameter-dependent geometry for the waveguide's metallic components and, in turn, in a parametrized model similar to

(2.7). However, in this case, only geometrical uncertainties are assumed. Considering the FE-based discretized model, the parametrized weak formulation reads

$$\text{Find } \mathbf{E}(\mathbf{y}) \in \mathbf{V} \text{ such that } a(\mathbf{y})\left(\mathbf{E}(\mathbf{y}), \mathbf{E}'\right) = \ell(\mathbf{E}'), \quad \forall \mathbf{E}' \in \mathbf{V}. \tag{6.21}$$

6.3.2 Numerical Results

As already mentioned, the input-to-output map is not smooth due to shifting resonance frequencies. Therefore, we expect that a polynomial surrogate model will offer a poor approximation accuracy. This is indeed the case, as shown in the cross-validation results presented in Figure 6.10c. The cross-validation error has been computed with a sample of size $Q_{\text{cv}} = 10^3$ for both a collocation and a gPC-based surrogate, the former using Algorithm 3.1 with Leja nodes and the latter using Algorithm 4.3 with $\kappa_{\max} = 10^2$. Both methods stagnate at an error well below the FEM accuracy of the discrete model. The approximation problems are particularly evident in the case of the collocation-based approximation, especially considering that the QoI, i.e. the reflection coefficient $|S_{11}|$, takes values around 0.82. A better accuracy is observed for the gPC-based surrogate, however, we cannot claim that the surrogate model can reliably replace the original model, at least if absolute accuracies of 10^{-2} or below are sought.

Next, we compare the moment estimation results of both polynomial-based approximation methods to the ones obtained with a MC sampling approach. Using

Table 6.3: Nominal parameter values for the resonant cavity filter.

Parameter	Nominal Value	Units
d_1	120	mm
d_2	60	mm
d_3	56	mm
d_{M}	56	mm
l_1	30	mm
l_2	30	mm
l_3	30	mm
l_{M}	60	mm
L_1	27	mm
L_2	18.5	mm
L_3	25	–

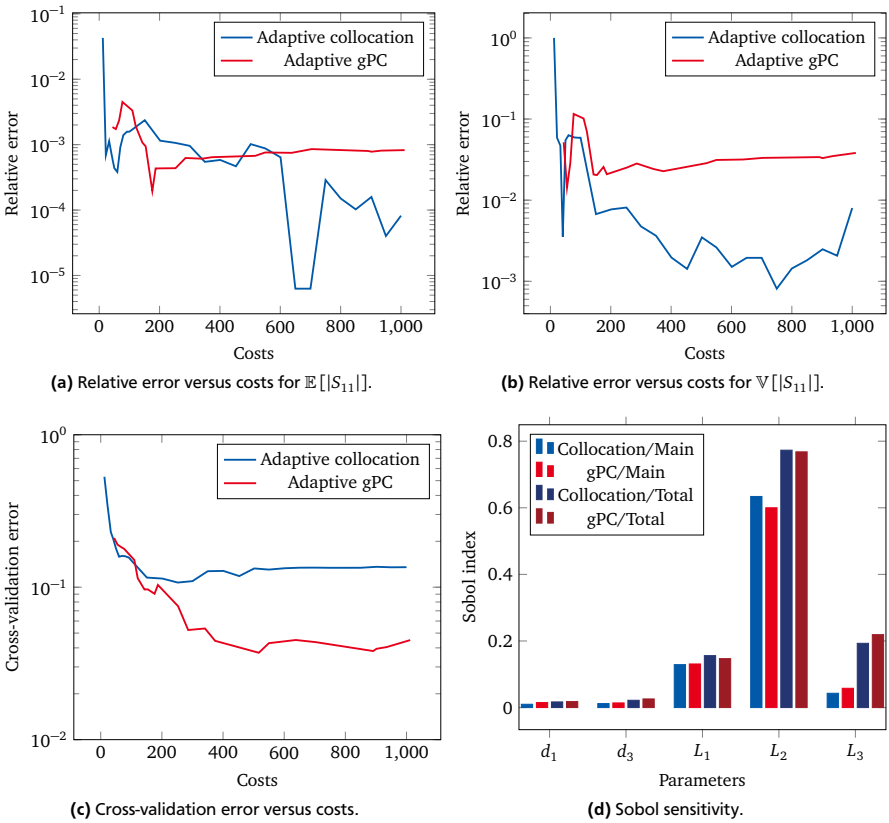

(a) Relative error versus costs for $\mathbb{E}[|S_{11}|]$.

(b) Relative error versus costs for $\mathbb{V}[|S_{11}|]$.

(c) Cross-validation error versus costs.

(d) Sobol sensitivity.

Figure 6.10: Uncertainty quantification and approximation results for the resonant cavity filter model. The approximation methods are the Leja-based dimension-adaptive stochastic collocation and the basis/sampling-adaptive gPC.

a sample of size $Q_{\mathrm{MC}} = 3 \cdot 10^3$, the MC-based expected value and variance, along with the corresponding RMSD and NRMSD, are

$$\begin{aligned}
&\mathbb{E}[|S_{11}|] = 0.8265,\\
&\mathbb{V}[|S_{11}|] = 0.0120,\\
&\mathrm{RMSD}[|S_{11}|] \approx 2 \cdot 10^{-3},\\
&\mathrm{NRMSD}[|S_{11}|] \approx 2.5 \cdot 10^{-3}.
\end{aligned}$$

The relative errors of the collocation and gPC-based moment estimations are given in Figures 6.10a and 6.10b, where the reference values correspond to the MC estimations. Although an absolute convergence cannot be observed for the maximum number of model evaluations, both methods yield relative errors close to or below the MC's NRMSD for significantly lower costs. For both moments, the gPC-based estimations seem to be more stable compared to the collocation-based ones. However, due to lack of absolute convergence, a general statement should be avoided. The results indicate that both approaches yield reliable moment estimations, despite their poor approximation results.

Finally, we use both surrogate models to perform a sensitivity analysis based on the Sobol method. For the gPC surrogate, the Sobol indices can be directly derived from the approximation's terms. Since the gPC variance estimations have been found to be reliable (see Figure 6.10b), we expect the gPC-based partial variances to be of similar accuracy, thus resulting in trustworthy Sobol index results. On the contrary, a sampling approach [135] must be used for the collocation-based surrogate, the approximation accuracy of which has been established to be very poor (see Figure 6.10c). Nevertheless, both methods yield similar results, as shown in Figure 6.10d. A possible explanation for this surprising result is that the surrogate models, although inaccurate in the approximation sense, have identified the input-output dependencies correctly. This hypothesis should be validated by a more thorough investigation.

In Figure 6.10d we plot both first and total order Sobol indices for all parameters with a contribution of 1% or above. The results indicate that the lengths of the 3 stubs are the most influential parameters, especially L_2. The parameters d_1 and d_3 corresponding to the irises of the first and third segments, respectively, are also found to have a non-negligible impact on $|S_{11}|$. The first order Sobol indices amount to approximately 83% of the total variance. The remaining 17% must be attributed to higher order interactions between the parameters. As can be observed in Figure 6.10d, the total Sobol indices of L_1 and L_3 are significantly higher than the first order ones. This is not observed for the remaining parameters, for which the first and total order Sobol indices are comparable. Hence, the non-first-order sensitivity can be attributed to an interaction between the parameters L_1 and L_3.

6.4 Concluding Remarks

In this section we have applied the methods presented in Chapters 3, 4, and 5 to three moderately high-dimensional problems, featuring 14, 14, and 11 uncertain parameters, respectively. The respective execution times for a single model evaluation are 2.5 μs, 60 s, and 40 minutes, considering the use of standard desktop machine. Using the analytical first model, we have again observed that the adaptive collocation and gPC approaches outperform significantly the adaptive tensor

approximations for moment estimations, this time in a higher-dimensional setting compared to Sections 3.6, 4.4, and 5.4. Then, the two computationally less expensive approaches have been tested on two real-world applications. For both considered applications, the performance of both methods can be regarded as comparable. A slight advantage can be given to the gPC approach in the case of not sufficiently smooth QoIs.

7 Conclusion and Outlook

In this final chapter, we present an overview of this thesis, as well as its most important conclusions. Moreover, we discuss a number of aspects which have not been considered in the context of this work and which should, in our opinion, be addressed in follow-up works.

7.1 Conclusion

In this work we have addressed the problem of high-dimensional UQ in the context of electromagnetic field (EMF) simulations using adaptive approximation methods. The dimension-adaptive stochastic collocation method and a self-developed adaptive gPC approach have been used for both approximation and moment estimation purposes. The latter has also been pursued with an approach based on tensor decompositions and adaptive cross approximations of tensors. The theory behind all examined approaches has been presented, along with their application in the UQ context.

Compared to classic, non-adaptive UQ approaches, all aforementioned adaptive methods have been found to yield significant advantages. The adaptive methods either reach similar errors at lower computational costs, or improve the accuracy for the same costs. This is achieved due to the parameter anisotropy which is inherent in all considered models. This anisotropy is exploited by the proposed methods in order to invest computational resources in an anisotropic way, thus focusing on parameters with significant impact on the QoI and neglecting insignificant parameters. In high-dimensional settings where non-adaptive methods become intractable, adaptive methods can enable UQ studies and yield accurate results.

The dimension-adaptive stochastic collocation has been tested for different choices of nested collocation points and input distributions. In each case, particular advantages and disadvantages are reported. A number of algorithms for the adaptive construction of gPC approximations have been proposed and compared against one another and against known adaptive approaches. The approach developed in this work is found to be superior to available adaptive approaches for the considered test case. The use of tensor decompositions in the UQ context is also demonstrated. Compared to full tensor approaches, tremendous complexity reductions are achieved, for similar accuracies.

Numerical experiments based on a fairly low dimensional model and a moderately high-dimensional model indicate that the stochastic collocation and adaptive gPC outperform the tensor-based approach by orders. In cases where the input-to-output map exhibits high regularity, the stochastic collocation method can be seen as the most advantageous option, particularly regarding surrogate model accuracy. However, as indicated by numerical experiments considering real-world

applications, this advantage vanishes for QoIs of lower regularity. In this case, the adaptive gPC can be regarded as either comparable or slightly advantageous, as it is less affected by the decreased regularity.

7.2 Outlook

We would like to mention some considerations, both connected to UQ as well as of a more general nature, which have not been addressed in this thesis and should be tackled in follow-up works.

1. A break-even between tensor decompositions, adaptive gPC and adaptive collocation remains to be found. Under strict theoretical assumptions, all methods are able to break the curse of dimensionality [38, 157]. However, such assumptions rarely hold in practical applications, where choosing the "best" method is often a problem-dependent task. Therefore, more investigations are needed in order to establish in which cases a method should be preferred over another.

2. Connected to the point above, another useful comparison for UQ practitioners would be between "on the fly" adaptive methods, as the ones presented in this work, and a priori dimension reduction approaches, e.g. active subspaces [43] or optimal sparse grids [10].

3. In the spirit of the multilevel Monte Carlo method [62, 152], multilevel or multifidelity collocation and gPC methods could be employed to further reduce computational costs. Such approaches have been studied in a mathematical context and for simple test cases [112, 151], however, their application to complicated engineering problems remains largely unexplored.

4. For all frequency-dependent models presented in this work, a single frequency has been considered. However, in most applications, the model's response over a frequency range is of interest. This poses an additional difficulty in terms of approximation and UQ, since the existence of poles and resonances in the frequency range will hinder the application of polynomial approximations. Therefore, efficient UQ approaches for the broadband case should be considered.

5. Connected to the point above, in this work we have always considered the case of a single scalar QoI. However, in many applications, multiple and/or vector-valued QoIs must be examined. How to effectively guide adaptivity in this setting remains a topic of active interest.

6. This work has considered a moderately high-dimensional UQ setting where the number of parameters is in the order of tens. Of particular interest, especially regarding real-world applications, are cases with hundreds or even thousands of parameters. Methods which allow parameter studies in such very high-dimensional settings should be examined.

7. As with the models employed in Sections 6.2 and 6.3, model symmetries are often exploited to reduce the computational cost of simulations. A particular difficulty in the UQ context is that random variations break symmetry. Realistic UQ studies should therefore employ full models, which correspond to increasing and often forbidding computational costs. How to effectively perform UQ in this regime is currently an active research topic and of particular interest in the case of periodic structures, e.g. metasurfaces [57, 129, 143].

8. Uncertainty quantification (UQ) is itself an interesting and challenging field, however, it becomes of actual practical interest in combination with other design tasks, e.g. optimization under uncertainty and robust optimization [77, 143]. The combination of already sophisticated techniques from different fields typically results in very challenging problems which are relevant in practical applications and should therefore be considered.

A Dielectric Slab Waveguide: Analytical Solution

We consider the dielectric slab waveguide problem presented in Section 2.1.1 and the non-parametric mathematical model (2.5). Writing the electric field $\mathbf{E}$ as the curl of a vector potential $\mathbf{A}$, such that $\mathbf{E} = \text{curl}\mathbf{A} = \nabla \times \mathbf{A}$, delivers the PDE

$$\nabla^2 \mathbf{A} + k^2 \mathbf{A} = \mathbf{0}. \tag{A.1}$$

Since we only assume TE modes propagating in the Cartesian z-direction, it is sufficient to consider only the z-component of the vector potential, such that $\mathbf{A}(x,y,z) = \mathbf{A}_z(x,y,x)\mathbf{e}_z$. Then, the solution of (A.1) takes the form

$$\begin{aligned}\mathbf{A}_z(x,y,z) = &\big(A\cos(k_x x) + B\sin(k_x x)\big)\\ &\big(C\cos(k_y y) + D\sin(k_y y)\big)\\ &\big(E\exp(-\jmath k_z z) + F\exp(\jmath k_z z)\big),\end{aligned} \tag{A.2}$$

with the separation equation

$$k_x^2 + k_y^2 + k_z^2 = \omega^2 \mu\varepsilon = k^2. \tag{A.3}$$

Because $\mathbf{E} = \nabla \times \mathbf{A}$, $\mathbf{H} = -\frac{1}{\jmath\omega\mu}\nabla \times \nabla \times \mathbf{A}$, and only the TE_{01} mode is taken into consideration, i.e. $k_x = 0$ and $k_y = \frac{\pi}{w}$, for $(x,y,z) \in \Omega$ ($\Omega = [0,w] \times [0,h] \times [0, 2d+l]$) holds that

$$E_x(x,y,z) = \begin{cases} E_0 \sin(k_x x)\big(\exp(-\jmath k_z z) + r_1 \exp(\jmath k_z z)\big), & z \in [0,d), \\ E_0 \sin(k_x x)\big(t_2 \exp(-\jmath k_{z1} z) + r_2 \exp(\jmath k_{z1} z)\big), & z \in [d, d+l), \\ E_0 \sin(k_x x) t_3 \exp(-\jmath k_z z), & z \in [d+l, 2d+l], \end{cases} \tag{A.4}$$

and

$$H_y(x,y,z) = \begin{cases} \frac{E_0 \sin(k_x x)}{Z_0}\big(\exp(-\jmath k_z z) - r_1 \exp(\jmath k_z z)\big), & z \in [0,d), \\ \frac{E_0 \sin(k_x x)}{Z_1}\big(t_2 \exp(-\jmath k_{z1} z) - r_2 \exp(\jmath k_{z1} z)\big), & z \in [d, d+l), \\ \frac{E_0 \sin(k_x x)}{Z_0} t_3 \exp(-\jmath k_z z), & z \in [d+l, 2d+l], \end{cases} \tag{A.5}$$

where k_z and k_{z1} are the propagation constants in vacuum, respectively, in the dielectric, $Z_0 = \frac{\omega\mu_0}{k_z}$ and $Z_1 = \frac{\omega\mu}{k_{z1}}$ are the corresponding wave impedances and r_1, r_2, t_2 and t_3 are coefficients modeling the reflection and transmission between the domains.

We apply the interface conditions at the boundary between the two materials, for $z = d$ and $z = d + l$, and use the continuity of the tangential field component of the magnetic and electric field to get four equations for the four unknowns r_1, r_2, t_2, and t_3, such that

$$\exp(-\jmath k_z d) + r_1 \exp(\jmath k_z d) = t_2 \exp(-\jmath k_{z1} d) + r_2 \exp(\jmath k_{z1} d), \tag{A.6a}$$

$$\frac{1}{Z_0}(\exp(-\jmath k_z d) - r_1 \exp(\jmath k_z d)) = \frac{1}{Z_1}(t_2 \exp(-\jmath k_{z1} d) - r_2 \exp(\jmath k_{z1} d)), \tag{A.6b}$$

$$t_2 \exp(-\jmath k_{z1}(d+l)) + r_2 \exp(\jmath k_{z1}(d+l)) = t_3 \exp(-\jmath k_{z1}(d+l)), \tag{A.6c}$$

$$\frac{1}{Z_1}(t_2 \exp(-\jmath k_{z1}(d+l)) - r_2 \exp(\jmath k_{z1}(d+l))) = \frac{1}{Z_0} t_3 \exp(-\jmath k_z (d+l)). \tag{A.6d}$$

The solutions of the system in (A.6) are

$$r_1 = \frac{1}{2}\exp(-\jmath k_z d)\left(t_2\left(1 - \frac{Z_0}{Z_1}\right)\exp(-\jmath k_{z1} d) + r_2\left(1 + \frac{Z_0}{Z_1}\right)\exp\left(\jmath k_{z_1} d\right)\right), \tag{A.7a}$$

$$r_2 = \frac{Z_1}{Z_1 - Z_0}\frac{2\exp(-\jmath k_z d)}{\exp\left(\jmath k_{z_1} d\right) - \exp(\jmath k_{z1}(2l+d))\left(\frac{Z_1+Z_0}{Z_1-Z_0}\right)^2}, \tag{A.7b}$$

$$t_2 = -\exp(2\jmath k_{z1}(l+d))\frac{Z_1 + Z_0}{Z_1 - Z_0} r_2, \tag{A.7c}$$

$$t_3 = \exp\left(\jmath k_{z_1}(d+l)\right)(t_2 \exp(-\jmath k_{z1}(d+l)) + r_2 \exp(\jmath k_{z1}(d+l))). \tag{A.7d}$$

Then, the reflection coefficient at port 1 is given by:

$$S_{11} = \frac{E_x^-(x,y,0)}{E_x^+(x,y,0)} = r_1. \tag{A.8}$$

List of Algorithms

List of Figures

List of Tables

List of Acronyms

ALS alternating least squares.

CC Clenshaw-Curtis.

CoV coefficient of variation.

CP canonical polyadic.

DMRG density matrix renormalization group.

DoF degrees of freedom.

EMF electromagnetic field.

FE finite element.

FEM finite element method.

GKP Gauss-Kronrod-Patterson.

gPC generalized polynomial chaos.

HF high frequency.

IGA isogeometric analysis.

LAR least angle regression.

LHS latin hypercube sampling.

LS least squares.

MC Monte Carlo.

MPS matrix product states.

NRMSD normalized root-mean-square deviation.

NURBS non-uniform rational B-splines.

PDE partial differential equation.

PDF probability density function.

PEC perfect electric conductor.

QoI quantity of interest.

RDP random discrete projection.

RHS right-hand side.

RMSD root-mean-square deviation.

RMSE root-mean-square error.

RV random variable.

SVD singular value decomposition.

TD total degree.

TE transverse electric.

TP tensor product.

TT tensor-train.

UQ uncertainty quantification.

References

[1] M. Ainsworth and J. T. Oden. *A posteriori error estimation in finite element analysis*. John Wiley & Sons, 2011.

[2] B. Arras, M. Bachmayr, and A. Cohen. Sequential sampling for optimal weighted least squares approximations in hierarchical spaces. *ArXiv e-prints*, 2018.

[3] O. Art and Z. Yi. Safe and effective importance sampling. *Journal of the American Statistical Association*, 95(449):135–143, 2000.

[4] R. Askey and J. Wilson. *Some Basic Hypergeometric Orthogonal Polynomials that Generalize Jacobi Polynomials*. Memoirs of the American Mathematical Society. AMS, 1985.

[5] A. C. M. Austin and C. D. Sarris. Efficient Analysis of Geometrical Uncertainty in the FDTD Method Using Polynomial Chaos With Application to Microwave Circuits. *IEEE Transactions on Microwave Theory and Techniques*, 61(12):4293–4301, 2013.

[6] I. Babuška, F. Nobile, and R. Tempone. A Stochastic Collocation Method for Elliptic Partial Differential Equations with Random Input Data. *SIAM Review*, 2:317–355, 2010.

[7] I. Babuška and W. C. Rheinboldt. A-posteriori error estimates for the finite element method. *International Journal for Numerical Methods in Engineering*, 12(10):1597–1615, 1978.

[8] I. Babuška, R. Tempone, and G. E. Zouraris. Galerkin Finite Element Approximations of Stochastic Elliptic Partial Differential Equations. *SIAM Journal on Numerical Analysis*, 42(2):800–825, 2004.

[9] J. Bäck, F. Nobile, L. Tamellini, and R. Tempone. Stochastic spectral Galerkin and collocation methods for PDEs with random coefficients: a numerical comparison. In *Spectral and High Order Methods for Partial Differential Equations*, volume 76 of *Lecture Notes in Computational Science and Engineering*, pages 43–62. Springer, 2011.

[10] J. Bäck, R. Tempone, F. Nobile, and L. Tamellini. On the optimal polynomial approximation of stochastic PDEs by Galerkin and collocation mehtods. *Mathematical Models and Methods in Applied Sciences*, 22(09), 2012.

[11] J. Ballani and L. Grasedyck. Hierarchical Tensor Approximation of Output Quantities of Parameter-Dependent PDEs. *SIAM/ASA Journal on Uncertainty Quantification*, 3(1):852–872, 2015.

[12] J. Ballani, L. Grasedyck, and M. Kluge. Black box approximation of tensors in hierarchical Tucker format. *Linear Algebra and its Applications*, 438(2):639 – 657, 2013.

[13] V. Barthelmann, E. Novak, and K. Ritter. High Dimensional Polynomial Interpolation on Sparse Grids. *Advances in Computational Mathematics*, 12(4):273–288, 2000.

[14] M. Baudin, A. Dutfoy, B. Iooss, and A.-L. Popelin. *OpenTURNS: An Industrial Software for Uncertainty Quantification in Simulation*, pages 2001–2038. Springer International Publishing, 2017.

[15] M. Bebendorf. Approximation of boundary element matrices. *Numerische Mathematik*, 86(4):565–589, 2000.

[16] M. Bebendorf. Adaptive cross approximation of multivariate functions. *Constructive Approximation*, 34(2):149–179, 2011.

[17] M. Bebendorf and S. Rjasanow. Adaptive low-rank approximation of collocation matrices. *Computing*, 70(1):1–24, 2003.

[18] R. E. Bellman. *Dynamic Programming*. Princeton University Press, 1957.

[19] P. Benner and J. Schneider. Uncertainty quantification for Maxwell's equations using stochastic collocation and model order reduction. *International Journal for Uncertainty Quantification*, 5(3), 2015.

[20] G. Beylkin and M. J. Mohlenkamp. Numerical operator calculus in higher dimensions. *Proceedings of the National Academy of Sciences*, 99(16):10246–10251, 2002.

[21] G. Beylkin and M. J. Mohlenkamp. Algorithms for numerical analysis in high dimensions. *SIAM Journal on Scientific Computing*, 26(6):2133–2159, 2005.

[22] G. Blatman and B. Sudret. An adaptive algorithm to build up sparse polynomial chaos expansions for stochastic finite element analysis. *Probabilistic Engineering Mechanics*, 25(2):183–197, 2010.

[23] G. Blatman and B. Sudret. Efficient computation of global sensitivity indices using sparse polynomial chaos expansions. 95:1216–1229, 11 2010.

[24] G. Blatman and B. Sudret. Adaptive Sparse Polynomial Chaos Expansion Based on Least Angle Regression. *Journal of Computational Physics*, 230(6):2345–2367, 2011.

[25] G. Blatman, B. Sudret, and M. Berveiller. Quasi random numbers in stochastic finite element analysis. *Mécanique & Industries*, 8(3):289–297, 2007.

[26] S. Börm and L. Grasedyck. Hybrid cross approximation of integral operators. *Numerische Mathematik*, 101(2):221–249, 2005.

[27] H.-J. Bungartz and M. Griebel. Sparse Grids. *Acta Numerica*, 13:147–269, 2004.

[28] J. Burkardt. The Truncated Normal Distribution, 2014. Department of Scientific Computing, Florida State University.

[29] T. Butler, C. Dawson, and T. Wildey. Propagation of uncertainties using improved surrogate models. *SIAM/ASA Journal on Uncertainty Quantification*, 1(1):164–191, 2013.

[30] G. T. Buzzard. Efficient Basis Change for Sparse-Grid Interpolating Polynomials with Application to T-Cell Sensitivity Analysis. *Computational Biology Journal*, 2013.

[31] R. E. Caflisch. Monte Carlo and quasi-Monte Carlo methods. *Acta Numerica*, 7:1–49, 1998.

[32] R. H. Cameron and W. T. Martin. Transformations of Wiener Integrals Under Translations. *Annals of Mathematics*, 45(2):386–396, 1944.

[33] J. D. Carroll and J.-J. Chang. Analysis of individual differences in multidimensional scaling via an n-way generalization of "Eckart-Young" decomposition. *Psychometrika*, 35(3):283–319, 1970.

[34] R. B. Cattell. "Parallel proportional profiles" and other principles for determining the choice of factors by rotation. *Psychometrika*, 9(4):267–283, 1944.

[35] P. Chen, A. Quarteroni, and G. Rozza. Reduced basis methods for uncertainty quantification. *SIAM/ASA Journal on Uncertainty Quantification*, 5(1):813–869, 2017.

[36] A. Chkifa, A. Cohen, G. Migliorati, F. Nobile, and R. Tempone. Discrete least squares polynomial approximation with random evaluations - application to parametric and stochastic elliptic PDEs. *ESAIM: Mathematical Modelling and Numerical Analysis*, 49(3):815–837, 2015.

[37] A. Chkifa, A. Cohen, and C. Schwab. High-Dimensional Adaptive Sparse Polynomial Interpolation and Applications to Parametric PDEs. *Foundations of Computational Mathematics*, 14(4):601–633, 2014.

[38] A. Chkifa, A. Cohen, and C. Schwab. Breaking the curse of dimensionality in sparse polynomial approximation of parametric PDEs. *Journal de Mathématiques Pures et Appliquées*, 103(2):400 – 428, 2015.

[39] C. W. Clenshaw and A. R. Curtis. A method for numerical integration on an automatic computer. *Numerische Mathematik*, 2(1):197–205, 1960.

[40] A. Cohen, M. A. Davenport, and D. Leviatan. On the stability and accuracy of least squares approximations. *Foundations of Computational Mathematics*, 13(5):819–834, 2013.

[41] A. Cohen and G. Migliorati. Optimal weighted least squares. *SMAI Journal of Computational Mathematics*, 3:181–203, 2017.

[42] A. Cohen and G. Migliorati. *Multivariate Approximation in Downward Closed Polynomial Spaces*, pages 233–282. Springer International Publishing, 2018.

[43] P. G. Constantine. *Active Subspaces: Emerging Ideas for Dimension Reduction in Parameter Studies*. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 2015.

[44] CST AG. CST STUDIO SUITE 2018.

[45] H. De Gersem, B. Masschaele, T. Roggen, E. Janssens, and T. T. Nguyen. Improved field post-processing for a Stern–Gerlach magnetic deflection magnet. *International Journal of Numerical Modelling: Electronic Networks, Devices and Fields*, 27(3):472–484, 2013.

[46] L. De Lathauwer, B. De Moor, and J. Vandewalle. A multilinear singular value decomposition. *SIAM Journal on Matrix Analysis and Applications*, 21(4):1253–1278, 2000.

[47] L. De Lathauwer, B. De Moor, and J. Vandewalle. On the Best Rank-1 and Rank-(R_1, R_2, ..., R_N) Approximation of Higher-Order Tensors. *SIAM Journal on Matrix Analysis and Applications*, 21(4):1324–1342, 2000.

[48] V. de Silva and L.-H. Lim. Tensor rank and the ill-posedness of the best low-rank approximation problem. *SIAM J. Matrix Anal. Appl.*, 30(3):1084–1127, 2008.

[49] B. J. Debusschere, H. N. Najm, P. P. Pébay, O. M. Knio, R. G. Ghanem, and O. P. Le Maître. Numerical Challenges in the Use of Polynomial Chaos Representations for Stochastic Processes. *SIAM Journal on Scientific Computing*, 26(2):698–719, 2004.

[50] S. Dolgov and D. Savostyanov. Alternating minimal energy methods for linear systems in higher dimensions. *SIAM Journal on Scientific Computing*, 36(5):A2248–A2271, 2014.

[51] C. Eckart and G. Young. The approximation of one matrix by another of lower rank. *Psychometrika*, 1(3):211–218, 1936.

[52] M. Espig, L. Grasedyck, and W. Hackbusch. Black box low tensor-rank approximation using fiber-crosses. *Constructive Approximation*, 30(3):557, 2009.

[53] J. Feinberg and H. P. Langtangen. Chaospy: An Open Source Tool for Designing Methods of Uncertainty Quantification. *Journal of Computational Science*, 11:46–57, 2015.

[54] L. Fejér. Mechanische Quadraturen mit Positiven Cotesschen Zahlen. *Mathematische Zeitschrift*, 37:287–309, 1933.

[55] S. Gabriel, R. W. Lau, and C. Gabriel. The dielectric properties of biological tissues: II. measurements in the frequency range 10 Hz to 20 GHz. *Physics in Medicine and Biology*, 41(11), 1996.

[56] S. Gabriel, R. W. Lau, and C. Gabriel. The dielectric properties of biological tissues: III. parametric models for the dielectric spectrum of tissues. *Physics in Medicine and Biology*, 41(11), 1996.

[57] N. Georg, D. Loukrezis, U. Römer, and S. Schöps. Uncertainty quantification for an optical grating coupler with an adjoint-based Leja adaptive collocation method. submitted.

[58] W. Gerlach and O. Stern. Das magnetische Moment des Silberatoms. *Zeitschrift für Physik A Hadrons und Nuclei*, 9(1):353–355, 1922.

[59] W. Gerlach and O. Stern. Der experimentell Nachweis der Richtungsquantelung im Magnetfeld. *Zeitschrift für Physik A Hadrons und Nuclei*, 9(1):349–352, 1922.

[60] T. Gerstner and M. Griebel. Dimension-Adaptive Tensor-Product Quadrature. *Computing*, 71(1):65–87, 2003.

[61] R. G. Ghanem and P. D. Spanos. *Stochastic Finite Elements: A Spectral Approach*. Springer-Verlag New York, Inc., New York, NY, USA, 1991.

[62] M. B. Giles. Multilevel Monte Carlo methods. *Acta Numerica*, 24:259–328, 2015.

[63] L. Giraldi, A. Litvinenko, D. Liu, H. G. Matthies, and A. Nouy. To Be or Not to Be Intrusive? The Solution of Parametric and Stochastic Equations - the "Plain Vanilla" Galerkin Case. *SIAM Journal on Scientific Computing*, 36(6), 2014.

[64] S. A. Goreinov, I. V. Oseledets, D. V. Savostyanov, E. E. Tyrtyshnikov, and N. L. Zamarashkin. How to find a good submatrix. *Research Report 08-10, ICM HKBU*, pages 08–10, 2008.

[65] S. A. Goreinov, E. E. Tyrtyshnikov, and N. L. Zamarashkin. A theory of pseudoskeleton approximations. *Linear Algebra and its Applications*, 261(1):1 – 21, 1997.

[66] L. Grasedyck. Hierarchical singular value decomposition of tensors. *SIAM Journal on Matrix Analysis and Applications*, 31(4):2029–2054, 2010.

[67] M. Griebel and J. Oettershagen. On tensor product approximation of analytic functions. *Journal of Approximation Theory*, 207:348 – 379, 2016.

[68] E. Griffiths. What is a model? 2010. URL: https:ncsu.edu/emily-griffiths/whatisamodel.pdf.

[69] W. Hackbusch. *Tensor Spaces and Numerical Tensor Calculus*. Springer-Verlag, Berlin, 2012.

[70] W. Hackbusch. Numerical Tensor Calculus. *Acta Numerica*, 23:651–742, 2014.

[71] W. Hackbusch and S. Kühn. A new scheme for the tensor representation. *Journal of Fourier Analysis and Applications*, 15(5):706–722, 2009.

[72] R. A. Harshman. Foundations of the PARAFAC procedure: Models and conditions for an exploratory multimodal factor analysis. *UCLA Working Papers in Phonetics*, 16:1–84, 1970.

[73] J. Håstad. Tensor rank is NP-complete. *Journal of Algorithms*, 11(4):644 – 654, 1990.

[74] J. S. Hesthaven, G. Rozza, and B. Stamm. *Certified Reduced Basis Methods for Parametrized Partial Differential Equations*. Springer-Verlag, Berlin, 2016.

[75] F. L. Hitchcock. The expression of a tensor or a polyadic as a sum of products. *Journal of Mathematics and Physics*, 6(1-4):164–189, 1927.

[76] T. J. Hughes, J. A. Cottrell, and Y. Bazilevs. Isogeometric Analysis: CAD, Finite Elements, NURBS, Exact Geometry and Mesh Refinement. *Computer Methods in Applied Mechanics and Engineering*, 194(39):4135–4195, 2005.

[77] I. G. Ion, Z. Bontinck, D. Loukrezis, U. Römer, O. Lass, S. Ulbrich, S. Schöps, and H. De Gersem. Robust shape optimization of electric devices based on deterministic optimization methods and finite-element analysis with affine parametrization and design elements. *Electrical Engineering*, 2018.

[78] J. D. Jakeman, M. S. Eldred, and K. Sargsyan. Enhancing ℓ_1-minimization estimates of polynomial chaos expansions using basis selection. *Journal of Computational Physics*, 289:18 – 34, 2015.

[79] R. Jankoski, U. Römer, and S. Schöps. Modeling of spatial uncertainties in the magnetic reluctivity. *COMPEL - The international journal for computation and mathematics in electrical and electronic engineering*, 36(4):1151–1167, 2017.

[80] P. Jantsch, C. G. Webster, and G. Zhang. On the Lebesgue constant of weighted Leja points for Lagrange interpolation on unbounded domains. *IMA Journal of Numerical Analysis*, 00:1–19, 2018.

[81] U. Jensen. Probabilistic risk analysis: Foundations and methods. *Journal of the American Statistical Association*, 97(459):925–925, 2002.

[82] J. Jin. *The Finite Element Method in Electromagnetics*. Wiley-IEEE Press, 3rd edition, 2014.

[83] A. Klimke and B. I. Wohlmuth. Algorithm 847: Spinterp: Piecewise Multilinear Hierarchical Sparse Grid Interpolation in MATLAB. *ACM Transactions on Mathematical Software*, 31(4):561–579, 2005.

[84] T. G. Kolda and B. W. Bader. Tensor decompositions and applications. *SIAM Review*, 51(3):455–500, 2009.

[85] D. A. Kopriva. *Spectral Approximation*, pages 3–38. Springer Netherlands, Dordrecht, 2009.

[86] A. S. Kronrod. *Nodes and weights of quadrature formulas. Sixteen-place tables*. Consultants Bureau, New York, 1965.

[87] O. P. Le Maître and O. M. Knio. *Spectral Methods for Uncertainty Quantification: With Applications to Computational Fluid Dynamics*. Scientific Computation. Springer Netherlands, 2010.

[88] O. P. Le Maître, O. M. Knio, H. N. Najm, and R. G. Ghanem. A Stochastic Projection Method for Fluid Flow: I. Basic Formulation. *Journal of Computational Physics*, 173(2):481–511, 2001.

[89] F. Leja. Sur certaines suites liées aux ensembles plans et leur application à la représentation conforme. *Annales Polonici Mathematici*, 4(1):8–13, 1957.

[90] C. Lemieux. *Monte Carlo and Quasi-Monte Carlo Sampling*. Springer Series in Statistics. Springer, Dordrecht, 2009.

[91] J. Li and D. Xiu. Evaluation of failure probability via surrogate models. *Journal of Computational Physics*, 229(23):8966–8980, 2010.

[92] W. K. Liu, T. Belytschko, and A. Mani. Probabilistic finite elements for nonlinear structural dynamics. *Computer Methods in Applied Mechanics and Engineering*, 56(1):61 – 81, 1986.

[93] W. K. Liu, T. Belytschko, and A. Mani. Random field finite elements. *International Journal for Numerical Methods in Engineering*, 23(10):1831–1845, 1986.

[94] W.-L. Loh. On latin hypercube sampling. *Annals of Statistics*, 24(5):2058–2080, 1996.

[95] D. Loukrezis. Dimension Adaptive Leja Interpolation (DALI). URL: https://github.com/dlouk/DALI. Accessed: 05.05.2018.

[96] D. Loukrezis, U. Römer, and H. De Gersem. Assessing the Performance of Leja and Clenshaw-Curtis Collocation for Computational Electromagnetics with Random Input Data. *ArXiv e-prints*, 2018.

[97] D. Loukrezis, U. Römer, T. Casper, S. Schöps, and H. De Gersem. High-dimensional uncertainty quantification for an electrothermal field problem using stochastic collocation on sparse grids and tensor train decompositions.

International Journal of Numerical Modelling: Electronic Networks, Devices and Fields, 31(2):e2222, 2017.

[98] F. Maradei, H. Ke, and T. H. Hubing. Full-wave model of frequency-dispersive media with Debye dispersion relation by circuit-oriented FEM. *IEEE Transactions on Electromagnetic Compatibility*, 51(2):312–319, 2009.

[99] S. Marelli and B. Sudret. UQLab User Manual - Polynomial Chaos Expansions, 2015.

[100] B. Masschaele, T. Roggen, H. De Gersem, E. Janssens, and T. T. Nguyen. Design of a Strong Gradient Magnet for the Deflection of Nanoclusters. *IEEE Transactions on Applied Superconductivity*, 22(3), 2012.

[101] H. G. Matthies and A. Keese. Galerkin Methods for Linear and Nonlinear Elliptic Stochastic Partial Differential Equations. *Informatik-Berichte der Technischen Universität Braunschweig*, 2003-08.

[102] G. Migliorati. Adaptive Polynomial Approximation by Means of Random Discrete Least Squares. In *ENUMATH*, volume 103 of *Lecture Notes in Computational Science and Engineering*, pages 547–554. Springer, 2013.

[103] G. Migliorati. Learning with discrete least squares on multivariate polynomial spaces using evaluations at random or low-discrepancy point sets. In *Machine Learning, Optimization, and Big Data*, pages 1–13. Springer International Publishing, 2015.

[104] G. Migliorati and F. Nobile. Analysis of discrete least squares on multivariate polynomial spaces with evaluations at low-discrepancy point sets. *Journal of Complexity*, 31(4):517–542, 2015.

[105] G. Migliorati, F. Nobile, E. von Schwerin, and R. Tempone. Approximation of Quantities of Interest in Stochastic PDEs by the Random Discrete L2 Projection on Polynomial Spaces. *SIAM Journal on Scientific Computing*, 35(3), 2013.

[106] G. Migliorati, F. Nobile, E. von Schwerin, and R. Tempone. Analysis of Discrete L2 Projection on Polynomial Spaces with Random Evaluations. *Foundations of Computational Mathematics*, 14(3):419–456, 2014.

[107] P. Monk. *Finite element methods for Maxwell's equations*. Oxford University Press, 2003.

[108] M. Motamed, F. Nobile, and R. Tempone. A Stochastic Collocation Method for the Second Order Wave Equation with a Discontinuous Random Speed. *Numerische Mathematik*, 123(3):493–536, 2013.

[109] S. Mustafa, A. M. Abbosh, and P. T. Nguyen. Modeling human head tissues using fourth-order Debye model in convolution-based three-dimensional finite-difference time-domain. *IEEE Transactions on Antennas and Propagation*, 62(3):1354–1361, 2014.

[110] A. Narayan and J. D. Jakeman. Adaptive Leja Sparse Grid Constructions for Stochastic Collocation and High-Dimensional Approximation. *SIAM Journal on Scientific Computing*, 36(6), 2014.

[111] J. C. Nédélec. Mixed finite elements in R^3. *Numerische Mathematik*, 35(3):315–341, 1980.

[112] L. W.-T. Ng and M. Eldred. Multifidelity Uncertainty Quantification Using Non-Intrusive Polynomial Chaos and Stochastic Collocation. In *53rd AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics and Materials Conference*. American Institute of Aeronautics and Astronautics, 2012.

[113] T. Nguyen, H. Mac, and S. Clénet. Uncertainty quantification using sparse approximation for models with a high number of parameters: Application to a magnetoelectric sensor. *IEEE Transactions on Magnetics*, 52:1–1, 2015.

[114] F. Nobile, L. Tamellini, and R. Tempone. *Comparison of Clenshaw–Curtis and Leja Quasi-Optimal Sparse Grids for the Approximation of Random PDEs*, pages 475–482. Springer International Publishing, 2015.

[115] F. Nobile and R. Tempone. Analysis and implementation issues for the numerical approximation of parabolic equations with random coefficients. *International journal for numerical methods in engineering*, 80(6-7):979–1006, 2009.

[116] F. Nobile, R. Tempone, and C. G. Webster. A Sparse Grid Stochastic Collocation Method for Partial Differential Equations with Random Input Data. *SIAM Journal on Numerical Analysis*, 46(5):2309–2345, 2008.

[117] F. Nobile, R. Tempone, and C. G. Webster. An Anisotropic Sparse Grid Stochastic Collocation Method for Partial Differential Equations with Random Input Data. *SIAM Journal on Numerical Analysis*, 46(5):2411–2442, 2008.

[118] A. Nouy. Proper generalized decompositions and separated representations for the numerical solution of high dimensional stochastic problems. *Archives of Computational Methods in Engineering*, 17(4):403–434, 2010.

[119] W. Oberkampf, T. Trucano, and C. Hirsch. Verification, validation, and predictive capability in computational engineering and physics. *Applied Mechanics Reviews*, 57(5):345–384, 2004.

[120] S. Orphanides. *Electromagnetic Waves and Antennas*. Rutgers University, 2016.

[121] I. Oseledets, S. Dolgov, V. Kazeev, T. Mach, O. Lebedeva, D. Savostyanov, P. Zhlobich, and L. Song. TT-Toolbox. URL: https://github.com/oseledets/TT-Toolbox/.

[122] I. Oseledets, D. Savostyanov, and E. Tyrtyshnikov. Tucker dimensionality reduction of three-dimensional arrays in linear time. *SIAM Journal on Matrix Analysis and Applications*, 30(3):939–956, 2008.

[123] I. Oseledets and E. Tyrtyshnikov. Breaking the curse of dimensionality, or how to use SVD in many dimensions. *SIAM Journal on Scientific Computing*, 31(5):3744–3759, 2009.

[124] I. Oseledets and E. Tyrtyshnikov. TT-cross approximation for multidimensional arrays. *Linear Algebra and its Applications*, 432(1):70 – 88, 2010.

[125] I. V. Oseledets. Tensor-Train Decomposition. *SIAM Journal on Scientific Computing*, 33(5):2295–2317, 2011.

[126] T. N. L. Patterson. The Optimum Addition of Points to Quadrature Formulae. *Mathematics of Computation*, 22(104):847–S31, 1968.

[127] A. Pels, Z. Bontinck, J. Corno, H. De Gersem, and S. Schöps. Optimization of a Stern-Gerlach Magnet by Magnetic Field-Circuit Coupling and Isogeometric Analysis. *IEEE Transactions on Magnetics*, 51(12), 2015.

[128] D. Perez-Garcia, F. Verstraete, M. M. Wolf, and J. I. Cirac. Matrix product state representations. *Quantum Information and Computation*, 7(5):401–430, 2007.

[129] A. Pitelet, E. Centeno, A. Moreau, N. Schmitt, C. Scheid, D. Loukrezis, and H. De Gersem. Influence of Spatial Dispersion on Surface Plasmons and Grating Couplers. submitted.

[130] M. M. Rao and R. J. Swift. *Probability Theory with Applications*. Springer-Verlag, 2006.

[131] M. E. Riley and R. V. Grandhi. Quantification of model-form and predictive uncertainty for multi-physics simulation. *Computers and Structures*, 89(11):1206 – 1213, 2011.

[132] U. Römer. *Numerical Approximation of the Magnetoquasistatic Model with Uncertainties and its Application to Magnet Design*. PhD thesis, Technische Universität Darmstadt, 2015.

[133] U. Römer, S. Schöps, and H. De Gersem. A Defect Corrected Finite Element Approach for the Accurate Evaluation of Magnetic Fields on Unstructured Grids. *Journal of Computational Physics*, 335:688–699, 2017.

[134] U. Römer, C. Schmidt, U. V. Rienen, and S. Schöps. Low-dimensional stochastic modeling of the electrical properties of biological tissues. *IEEE Transactions on Magnetics*, 53(6):1–4, 2017.

[135] A. Saltelli. Making best use of model evaluations to compute sensitivity indices. *Computer Physics Communications*, 145(2):280–297, 2002.

[136] A. Saltelli. Sensitivity analysis for importance assessment. *Risk Analysis*, 22(3):579–590, 2002.

[137] D. V. Savostyanov. Quasioptimality of maximum-volume cross interpolation of tensors. *Linear Algebra and its Applications*, 458:217 – 244, 2014.

[138] D. V. Savostyanov and I. V. Oseledets. Fast adaptive interpolation of multi-dimensional arrays in tensor train format. In *The 2011 International Workshop on Multidimensional (nD) Systems*, pages 1–8. IEEE, 2011.

[139] B. Schieche. *Unsteady Adaptive Stochastic Collocation on Sparse Grids*. PhD thesis, Technische Universität Darmstadt, 2012.

[140] C. Schillings and C. Schwab. Sparse, adaptive Smolyak quadratures for Bayesian inverse problems. *Inverse Problems*, 29(6), 2013.

[141] C. Schmidt, P. Grant, M. Lowery, and U. van Rienen. Influence of uncertainties in the material properties of brain tissue on the probabilistic volume of tissue activated. *IEEE Transactions on Biomedical Engineering*, 60(5):1378–1387, 2013.

[142] E. Schmidt. Zur Theorie der linearen und nichtlinearen Integralgleichungen. *Mathematische Annalen*, 63(4):433–476, 1907.

[143] N. Schmitt, N. Georg, G. Brière, D. Loukrezis, S. Héron, S. Lanteri, U. Römer, H. De Gersem, and P. Genevet. Optimization and Uncertainty Quantification of Gradient Index Metasurfaces. submitted.

[144] U. Schollwöck. The density-matrix renormalization group in the age of matrix product states. *Annals of Physics*, 326(1):96 – 192, 2011.

[145] S. A. Smolyak. Quadrature and Interpolation Formulas for Tensor Products of Certain Classes of Functions. *Soviet Mathematics Doklady*, 4:240–243, 1963.

[146] I. M. Sobol. Global Sensitivity Indices for Nonlinear Mathematical Models and their Monte Carlo Estimates. *Mathematics and Computers in Simulation*, 55(1):271–280, 2001.

[147] P. Sochala and O. L. Maitre. Polynomial chaos expansion for subsurface flows with uncertain soil parameters. *Advances in Water Resources*, 62:139–154, 2013.

[148] A. Sommariva. Fast Construction of Fejér and Clenshaw–Curtis Rules for General Weight Functions. *Computers & Mathematics with Applications*, 65(4):682–693, 2013.

[149] B. Sudret. Global sensitivity analysis using polynomial chaos expansions. *Reliability Engineering & System Safety*, 93(7):964 – 979, 2008.

[150] L. Tamellini and F. Nobile. Sparse Grids MATLAB Kit. URL: http://csqi.epfl.ch. Accessed: 15.08.2017.

[151] A. L. Teckentrup, P. Jantsch, C. Webster, and M. Gunzburger. A multilevel stochastic collocation method for partial differential equations with random input data. *SIAM/ASA Journal on Uncertainty Quantification*, 3(1):1046–1074, 2015.

[152] A. L. Teckentrup, R. Scheichl, M. B. Giles, and E. Ullmann. Further analysis of multilevel Monte Carlo methods for elliptic PDEs with random coefficients. *Numerische Mathematik*, 125(3):569–600, 2013.

[153] L. N. Trefethen. Is Gauss Quadrature Better than Clenshaw–Curtis? *SIAM Review*, 50(1):67–87, 2008.

[154] L. N. Trefethen. Six myths of polynomial interpolation and quadrature. *Mathematics Today*, 2011.

[155] L. R. Tucker. Some mathematical notes on three-mode factor analysis. *Psychometrika*, 31(3):279–311, 1966.

[156] E. Tyrtyshnikov. Incomplete cross approximation in the mosaic-skeleton method. *Computing*, 64(4):367–380, 2000.

[157] A. Uschmajew. *Zur Theorie der Niedrigrangapproximation in Tensorprodukten von Hilberträumen*. PhD thesis, Technische Universität Berlin, 2013.

[158] R. Verfürth. *A Posteriori Error Estimation Techniques for Finite Element Methods*. Oxford University Press, 2013.

[159] E. Waring. Problems concerning interpolations. *Philosophical Transactions of the Royal Society of London*, 69:59–67, 1779.

[160] S. R. White. Density matrix formulation for quantum renormalization groups. *Physical Review Letters*, 69:2863–2866, 1992.

[161] N. Wiener. The Homogeneous Chaos. *American Journal of Mathematics*, 60(4):897–936, 1938.

[162] D. Xiu. *Numerical Methods for Stochastic Computations: A Spectral Method Approach*. Princeton University Press, Princeton, NJ, USA, 2010.

[163] D. Xiu and J. S. Hesthaven. High-Order Collocation Methods for Differential Equations with Random Inputs. *SIAM Journal on Scientific Computing*, 27(3):1118–1139, 2005.

[164] D. Xiu and G. E. Karniadakis. The Wiener-Askey Polynomial Chaos for Stochastic Differential Equations. *SIAM Journal on Scientific Computing*, 24(2):619–644, 2002.

[165] D. Xiu and G. E. Karniadakis. Supersensitivity due to uncertain boundary conditions. *International Journal for Numerical Methods in Engineering*, 61(12):2114–2138, 2004.

[166] S. Zaglmayr. *High order finite element methods for electromagnetic field computation*. PhD thesis, Johannes Kepler Universität Linz, 2006.

[167] Z. Zhang, X. Yang, I. V. Oseledets, G. E. Karniadakis, and L. Daniel. Enabling High-Dimensional Hierarchical Uncertainty Quantification by ANOVA and Tensor-Train Decomposition. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 34(1):63–76, 2015.

Acknowledgements

I would like to express my sincerest gratitude to:

- Prof. Dr.-Ing. Herbert De Gersem, for the invitation[5] to work at the Institut für Theorie Elektromagnetischer Felder (TEMF) as a research assistant, the opportunity to pursue a thesis topic of my choice, the supervision of my teaching and research activities, and for always having his door open whenever I needed his advice.
- Prof. Dr.-Ing. Ulrich Römer, for more than 3 years of mentoring and day-to-day supervision, either as a group leader at TEMF, or as a professor at the Technische Universität Braunschweig.
- Prof. Dr. rer. nat. Sebastian Schöps, for our numerous, scientific and non-scientific, discussions.
- Dr.-Ing. Stefan Reitzinger, Thorben Casper, Andreas Pels, Dr.-Ing. Jacopo Corno, Zeger Bontinck, Dr.-Ing. Nikolai Schmitt, Armel Pitelet, and Dr.-Ing. Nicolas Marsic, for providing me computational models to play around with various UQ methods.
- Armin Galetzka and Ion Gabriel Ion, for their research work under my supervision, either as thesis students or as student research assistants, as well as for proofreading some parts of this thesis.
- My colleagues at TEMF, for our daily discussions, jokes, and arguments, especially Radoslav Jankoski and Niklas Georg for being excellent office-mates, as well as Yvonne Späck-Leigsnering and Yun Ouédraogo for their comments on this thesis.
- My parents, Kostantinos and Maria, my sister, Irini, and my grandmother, Irini, as well as the non-nucleus members of our quite extended family, for their unconditional love, support, and encouragement.
- All of my friends, for being my friends. No names here, you know who you are.
- The “long halls of science and all the lunatics committed there”[6].

This work has been partially supported by the “Excellence Initiative” of the German Federal and State Governments and the Graduate School of Computational Engineering at the Technische Universität Darmstadt.

[5] Herbert, as they say, third time’s a charm.

[6] Clutch, *10001110101*.

Academic Career

Education

Aug. 2015 – Feb. 2019 Doctorate in Electrical Engineering and Information Technology, Technische Universität (TU) Darmstadt, Darmstadt, Germany.

Oct. 2012 – May 2014 Master of Science in Simulation Sciences, Rheinisch - Westfälische Technische Hochschule (RWTH) Aachen, Aachen, Germany.

Oct. 2005 – Jul. 2012 Diploma of Advanced Studies in Electrical and Computer Engineering, National Technical University of Athens (NTUA), Athens, Greece.

Awards

Nov. 2016 Athene-Preis für Gute Lehre (Athena Award for Good Teaching), TU Darmstadt.

Sep. 2015 Springorum Denkmüntze (Springorium Commemorative Coin), awarded for excellence in M.Sc. studies, RWTH Aachen.

Certifications

Ongoing Zertifikat Hochschullehre (Certificate for University Teaching), TU Darmstadt.

Publications

1. D. Loukrezis, U. Römer, and H. De Gersem, *Assessing the Performance of Leja and Clenshaw-Curtis Collocation for Computational Electromagnetics with Random Input Data*, International Journal for Uncertainty Quantification, 2019 (accepted).

2. N. Schmitt, N. Georg, G. Briére, D. Loukrezis, S. Héron, S. Lanteri, C. Klitis, M. Sorel, U. Römer, H. De Gersem, S. Vézian, and P. Genevet, *Optimization and uncertainty quantification of gradient index metasurfaces*, Optical Materials Express, 2019.

3. I. G. Ion, Z. Bontinck, D. Loukrezis, U. Römer, O. Lass, S. Schöps, S. Ulbrich, and H. De Gersem, *Robust Shape Optimization of Electric Devices Based on Deterministic Optimization Methods and Finite Element Analysis With Affine Decomposition and Design Elements*, Electrical Engineering, ISSN 1432-0487, 2018.

4. D. Loukrezis, U. Römer, T. Casper, S. Schöps, and H. De Gersem, *High-Dimensional Uncertainty Quantification for an Electrothermal Field Problem Using Stochastic Collocation on Sparse Grids and Tensor Train Decompositions*, International Journal of Numerical Modeling: Electronic Networks, Devices and Fields, 31 (2), ISSN 1099-1204, 2018.

Preprints

1. N. Georg, D. Loukrezis, U. Römer, and S. Schöps, *Uncertainty Quantification for an Optical Grating Coupler with an Adjoint Error-Based Leja Adaptive Collocation Method*, arXiv e-print.

2. A. Pitelet, E. Centeno, A. Moreau, N. Schmitt, C. Scheid, D. Loukrezis, and H. De Gersem, *Influence of Spatial Dispersion on Surface Plasmons and Grating Couplers*, submitted.